arXiv:1401.2436v1 [cs.CC] 10 Jan 2014

Hardness of Robust Graph Isomorphism, Lasserre Gaps, and Asymmetry of Random Graphs

Ryan O'Donnell*
Department of Computer Science
Carnegie Mellon University
odonnell@cs.cmu.edu

John Wright*†
Department of Computer Science
Carnegie Mellon University
jswright@cs.cmu.edu

Chenggang Wu‡
IIIS
Tsinghua University
wcg06@mails.tsinghua.edu.cn

Yuan Zhou*†§
Department of Computer Science
Carnegie Mellon University
yuanzhou@cs.cmu.edu

January 13, 2014

Abstract

Building on work of Cai, Fürer, and Immerman [CFI92], we show two hardness results for the Graph Isomorphism problem. First, we show that there are pairs of nonisomorphic n-vertex graphs G and H such that any sum-of-squares (SOS) proof of nonisomorphism requires degree $\Omega(n)$. In other words, we show an $\Omega(n)$-round integrality gap for the Lasserre SDP relaxation. In fact, we show this for pairs G and H which are not even $(1 - 10^{-14})$-isomorphic. (Here we say that two n-vertex, m-edge graphs G and H are α-isomorphic if there is a bijection between their vertices which preserves at least αm edges.) Our second result is that under the R3XOR Hypothesis [Fei02] (and also any of a class of hypotheses which generalize the R3XOR Hypothesis), the *robust* Graph Isomorphism problem is hard. I.e. for every $\epsilon > 0$, there is no efficient algorithm which can distinguish graph pairs which are $(1-\epsilon)$-isomorphic from pairs which are not even $(1-\epsilon_0)$-isomorphic for some universal constant ϵ_0. Along the way we prove a robust asymmetry result for random graphs and hypergraphs which may be of independent interest.

*Supported by NSF grants CCF-0747250 and CCF-1116594 and a grant from the MSR–CMU Center for Computational Thinking.
†Some of this research done with visiting Toyota Technological Institute at Chicago.
‡This work was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61061130540. Some of this research done while visiting Carnegie Mellon University.
§Also supported by a grant from the Simons Foundation (Award Number 252545).

1 Introduction

The GRAPHISOMORPHISM problem is one of the most intriguing and notorious problems in computational complexity theory (we will also refer to it as GISO for short); we refer to [KST93, Bab95, AK05, Köb06, Cod11] for surveys. It was famously referred to as a "disease" over 35 years ago [RC77] and maintains its infectious status to this day. Together with FACTORING, it is one of the very rare problems in NP which is not known to be in P but which is believed to not be NP-hard [Bab85, BHZ87, Sch88] (according to standard complexity-theoretic assumptions). Both problems also admit an algorithm with running time "subexponential" (or "moderately exponential") in the natural witness size. In the case of GRAPHISOMORPHISM on n-vertex graphs, the natural witness size is $\log_2(n!) = \Theta(n \log n)$, but the best known algorithm due to Luks solves the problem in time $2^{O(\sqrt{n \log n})}$ [BL83].

In the same breath we might mention the problems GAPSVP$_{\sqrt{n}}$ (approximating the shortest vector in an n-dimensional lattice to factor $\sqrt{n}$) and UNIQUEGAMES$_\epsilon$ (the Unique Games problem proposed by Khot [Kho02]). The former is not NP-hard subject to standard complexity-theoretic assumptions [GG00, AR05], though we don't know any subexponential-time algorithm. The latter has a subexponential-time algorithm [ABS10]; whether it is NP-hard or in P (or neither) is hotly contested. The potential hardness of FACTORING and GAPSVP — even under certain average-case distributions — is well enough entrenched that many cryptographic protocols are based on it. (The same is true of random 3XOR with noise, more on which later.) On the other hand, for GRAPHISOMORPHISM and UNIQUEGAMES$_\epsilon$ we do not know any way of generating "hard-seeming instances"; indeed, some experts have speculated that GRAPHISOMORPHISM may be in P, or at least have a $2^{\mathrm{polylog}(n)}$-time algorithm.

In this paper we investigate hardness results for the GISO problem. Since GISO may well be in P, let us discuss what this may mean. One direction would be to show that GISO is hard for small complexity classes. This has been pursued most successfully by Torán [Tor04], who has shown that GISO is hard for the class DET. This is essentially the class of problems equivalent to computing the determinant; it contains #L and is contained in TC^1. It is not known whether GISO is P-hard.

In this paper, however, we are concerned with hardness results well above P. We have two related main results. The first is that solving GISO via the Lasserre/SOS hierarchy requires $2^{\Omega(n)}$ time (i.e., $\Omega(n)$ rounds/degree). This generalizes the result of Cai, Fürer, and Immerman [CFI92] showing that the frequently effective $o(n)$-dimensional Weisfeiler–Lehman algorithm fails to solve GISO; it also gives even more evidence that any subexponential-time algorithm for GISO requires algebraic, non-local techniques.

Our second result is concerned with the problem of *robust* graph isomorphism, ROBUSTGISO. Roughly speaking, this is the following problem: given two graphs which are almost isomorphic, find an "almost-isomorphism". ROBUSTGISO is strictly harder than GISO and the fact that it concerns "isomorphisms with noise" seems to rule out all algebraic techniques. We show that ROBUSTGISO is at least as hard as random 3XOR with noise; hence ROBUSTGISO has no polynomial-time algorithm assuming the well-known R3XOR Hypothesis of Feige [Fei02]. In fact, it's possible that the R3XOR problem requires $2^{n^{1-o(1)}}$ time, which would mean that ROBUSTGISO is much harder than GISO itself.

1.1 SOS/Lasserre gaps

The most well-known heuristic for GRAPHISOMORPHISM (and the basis of most practical algorithms — e.g., "nauty" [McK81]) is the Weisfeiler–Lehman (WL) algorithm [Wei76] and its "higher dimensional" generalizations. To describe the basic algorithm we need the notion of a colored graph. This is simply a graph, together with a function mapping the vertices to a finite set of colors; equivalently, a graph with its vertices partitioned into "color classes". Isomorphisms involving colored graphs are always assumed to preserve colors. Let G be a colored graph on the n-vertex set V. A *color refinement step* refers to the following procedure: for each $v \in V$, one determines the multiset C_v of colors in the neighborhood of v; then one recolors each v with color C_v. Now the basic WL algorithm, when given graphs G and H, repeatedly applies refinement to each of them until the colorings stabilize. (Initially, the graphs are treated as having just one color class.) At the end, if G and H have the same number of vertices of each color the WL algorithm outputs that they are "maybe isomorphic"; otherwise, it (correctly) outputs that they are "not isomorphic".

Note that after the initial refinement step, a graph's vertices are colored according to their degree. Thus two d-regular graphs are always reported as "maybe isomorphic" by the basic WL algorithm. On the other hand,

the heuristic is powerful enough to work correctly for all trees and for almost all n-vertex graphs in the Erdős–Rényi $\mathcal{G}(n, 1/2)$ model [BES80, BK79]. (We say the heuristic "works correctly" on a graph G if the stabilized coloring for G is distinct from the stabilized coloring of any graph not isomorphic to G.) To overcome WL's failure for regular graphs, several researchers independently introduced the "k-dimensional generalization" WL^k (see [Wei76, CFI92] for discussion). Briefly, in the WL^k heuristic, each k-tuple of vertices has a color, and color refinement involves looking at all "neighbors" of each k-tuple of vertices $(v_1, \ldots, v_k)$ (where the neighbors are all tuples of the form $(v_1, \ldots, v_{i-1}, u_i, v_{i+1}, \ldots, v_k)$, where $\{u_i, v_i\}$ is an edge). The WL^k heuristic can be performed in time $n^{k+O(1)}$ and is thus a polynomial-time algorithm for any constant k.

The WL^k heuristic is very powerful. For example, it is known to work correctly in polynomial time for all graphs which exclude a fixed minor [Gro12], a class which includes all graphs of bounded treewidth or bounded genus. Spielman's $2^{\widetilde{O}(n^{1/3})}$-time graph isomorphism algorithm [Spi96] for strongly regular graphs is achieved by WL^k with $k = \widetilde{O}(n^{1/3})$. The WL^k algorithm with $k = O(\sqrt{n})$ is also a key component in the $2^{O(\sqrt{n \log n})}$-time GI$_{\text{SO}}$ algorithm [BL83]. Throughout the '80s there was some speculation that GI$_{\text{SO}}$ might be solvable on all graphs by running the WL^k algorithm with $k = O(\log n)$ or even $k = O(1)$. However this was disproved in the notable work of Cai, Fürer, and Immerman [CFI92], which showed the existence of nonisomorphic n-vertex graphs G and H which are not distinguished by WL^k unless $k = \Omega(n)$.[1]

The GRAPHISOMORPHISM problem can be thought of as kind of constraint satisfaction problem (CSP), and readers familiar with LP/SDP hierarchies for CSPs might see an analogy between k-dimensional WL and level-k LP/SDP relaxations. A very interesting recent work of Atserias and Maneva [AM13] (see also [GO12]) shows that this is more than just an analogy — it shows that the power of WL^k is precisely sandwiched between that of the kth and $(k+1)$st level of the canonical *Sherali–Adams LP hierarchy* [SA90]. (In fact, it had long been known [RSU94] that WL^1 is equivalent in power to the basic LP relaxation of GI$_{\text{SO}}$.) This gives a very satisfactory connection between standard techniques in optimization algorithms and the best known non-algebraic/local techniques for GI$_{\text{SO}}$.

This connection raises the question of whether stronger LP/SDP hierarchies might prove more powerful than WL^k in the context of GI$_{\text{SO}}$. The strongest such hierarchy known is the "SOS (sum-of-squares) hierarchy" due to Lasserre [Las00] and Parrilo [Par00]. Very recent work [BBH$^+$12, OZ13, KOTZ13] in the field of CSP approximability has shown that $O(1)$ levels of the SOS hierarchy can succeed where $\omega(1)$ levels of weaker SDP hierarchies fail; in particular, this holds for the hardest known instances of UNIQUEGAMES$_\epsilon$ [BBH$^+$12]. This raises the question of whether there might be a subexponential-time algorithm based SOS which solves GRAPHISOMORPHISM.

We answer this question negatively. Our first main result is that a variant of the Cai–Fürer–Immerman instances also fools $\Omega(n)$ levels of the SOS hierarchy. In fact, we achieve a "constant factor Lasserre gap with perfect completeness". To explain this we require a definition:

Definition 1. *Let G and H be nonempty n-vertex graphs. For $0 \leq \beta \leq 1$, we say that a bijection $\pi \colon V(G) \to V(H)$ is an* α-isomorphism *if*

$$\frac{|\{(u,v) \in E(G) : (\pi(u), \pi(v)) \in E(H)\}|}{\max\{|E(G)|, |E(H)|\}} \geq \alpha.$$

In this case we say that G and H are α-isomorphic.

Observe that this definition is symmetric in G and H. The two graphs are isomorphic if and only if they are 1-isomorphic. We will almost always consider the case that G and H have the same number of edges. We prove:

Theorem 1.1. *For infinitely many n, there exist pairs of n-vertex, $O(n)$-edge graphs G and H such that:*

- *G and H are not $(1 - 10^{-14})$-isomorphic;*
- *any SOS refutation of the statement "G and H are isomorphic"[2] requires degree $\Omega(n)$.*

[1] Actually, G and H are colored graphs in [CFI92]'s construction, with each color class having size at most 4. It is often stated that the colors can be replaced by gadgets while keeping the number of vertices $O(n)$. We do not find this to be immediately obvious. However it does follow from the asymmetry of random graphs, as we will see in this paper.

[2] When naturally encoded.

A word on our techniques. The essence of the Cai–Fürer–Immerman construction is to take a 3-regular expander graph and replace each vertex by a certain 10-vertex gadget (originally appearing in [Imm81] and also sometimes called a "Fürer gadget"). This gadget is closely related to 3-variable equations modulo 2 (as observed by several authors, e.g. [Tor04]); indeed, it may be described as the "label-extended graph" of the 3XOR constraint. The reader may therefore recognize the [CFI92] WL^k lower bound as stemming from the difficulty of refuting unsatisfiable, expanding 3XOR CSP instances by "local" means. This should make our Theorem 1.1 look plausible in light of the Grigoriev [Gri01] and Schoenebeck [Sch08] SOS/Lasserre lower bounds.

Nevertheless, obtaining Theorem 1.1 is not automatic. For one, we still lack a complete theory of reductions within the SOS hierarchy (though see [Tul09]). Second, the pair of graphs constructed by [CFI92] only differ by one edge. More tricky is the issue of removing the "colors" from the [CFI92] construction. We do not see an easy gadget-based way of doing this without sacrificing on the $\Omega(n)$ degree. To handle this we have to: a) modify the [CFI92] construction somewhat to make the two graphs differ by a constant fraction; b) prove that random (hyper)graphs are "robustly asymmetric" — i.e., "far from having nontrivial automorphisms"; c) use the robust asymmetry property to remove the "colors". The result in b), described below in Section 1.3, qualitatively generalizes work of Erdős and Rényi [ER63] and may be of independent interest.

In an independent work, Codenotti, Schoenebeck, and Snook [CSS14] have shown a conclusion similar to our Theorem 1.1. Their main result is that there are expander graphs G and H which are not isomorphic, but any SOS refutation of the statement "G and H are isomorphic" requires degree $\Omega(n)$. As in our work, their proof combines the Cai–Fürer–Immerman construction with the Schoenebeck [Sch08] SOS/Lasserre lower bounds.

1.2 Robust graph isomorphism

Our second main result concerns the RobustGiso problem. To introduce it, we might imagine an algorithm trying to recover an isomorphism between G and H, where H is formed by permuting the vertices of G and then introducing a small amount of *noise* — say, adding and deleting an ϵ fraction of edges. Thinking of Giso as a CSP, we are concerned with finding "almost-satisfying" solutions on "almost-satisfiable" instances. For example, suppose we are given graphs G and H which are promised to be $(1-\epsilon)$-isomorphic. Can we efficiently find a $(1-2\epsilon)$-isomorphism? A $(1-\sqrt{\epsilon})$-isomorphism? A $(1-\frac{1}{\log(1/\epsilon)})$-isomorphism?

Definition 2. *We say an algorithm $\mathcal{A}$ solves the* RobustGiso *problem if there is a function $r\colon [0,1] \to [0,1]$ satisfying $r(\epsilon) \to 0$ as $\epsilon \to 0^+$ such that whenever $\mathcal{A}$ is given any pair of graphs which are $(1-\epsilon)$-isomorphic, $\mathcal{A}$ outputs a $(1-r(\epsilon))$-isomorphism between them.*[3]

Remark 1.1. *In particular, $\mathcal{A}$ must solve the* Giso *problem, because given isomorphic graphs with at most m edges we can always take $\epsilon > 0$ small enough so that $r(\epsilon) < 1/m$.*

The analogous problem of robust satisfaction algorithms for CSPs over constant-size domains was introduced by Zwick [Zwi98] and has proved to be very interesting. Guruswami and Zhou [GZ11] conjectured that the CSPs which have efficient robust algorithms (subject to $\mathsf{P} \neq \mathsf{NP}$) are precisely those of "bounded width" — roughly speaking, those that do not encode equations over abelian groups. This conjecture was recently confirmed by Barto and Kozik [BK12], following partial progress in [KOT+12, DK12]. The problem of RobustGiso was first introduced explicitly by Wu et al. [WYZV13]; their results included, among another things, a nontrivial efficient RobustGiso algorithm for trees. It is an interesting question to find efficient RobustGiso algorithms for more classes of graphs on which Giso is solvable — e.g., planar graphs, or bounded treewidth graphs.

The graph isomorphism seems to share some of the flavor of "unbounded width" CSPs such as 3XOR; these CSPs have the property that special algebraic methods (namely, Gaussian elimination) are available on satisfiable instances, but these methods break down once there is a small amount of noise. Indeed, the $2^{O(\sqrt{n \log n})}$-time algorithm for Giso is a somewhat peculiar mix of group theory and "local methods" (namely, Weisfeiler–Lehman). Generalizing from Giso to RobustGiso seems like it might rule out applicability of group-theoretic methods, thereby making the problem much harder. Our second main theorem in a sense confirms this. Roughly speaking, it shows that RobustGiso is hard assuming it is hard to distinguish random 3XOR instances from random instances with a planted solution and slight noise:

[3] We could also consider the easier task of distinguishing pairs which are $(1-\epsilon)$-isomorphic from pairs which are not $(1-r(\epsilon))$-isomorphic. In fact, our hardness result will hold even for this easier problem.

Theorem 1.2. *Assume Feige's* R3XOR *Hypothesis [Fei02]. Then there is no polynomial-time algorithm for* RobustGiso. *More precisely, there exists $\epsilon_0 > 0$, such that suppose there exists $\epsilon > 0$ and a $t(n)$-time algorithm which can distinguish $(1-\epsilon)$-isomorphic n-vertex, m-edge graph pairs from pairs which are not even $(1-\epsilon_0)$-isomorphic (where $m = O(n)$). Then there is a universal constant $\Delta \in \mathbb{Z}^+$ and a $t(O(n))$-time algorithm which outputs "typical" for almost all n-variable, Δn-constraint instances of the* 3XOR *problem, yet which never outputs "typical" on instances which are $(1-\Theta(\epsilon))$-satisfiable.*

Here we refer to:

Feige's R3XOR **Hypothesis.** *For every fixed $\epsilon > 0$, $\Delta \in Z^+$, there is no polynomial time algorithm which on almost all* 3XOR *instances with n variables and $m = \Delta n$ constraints outputs "typical", but which never outputs "typical" on instances which an assignment satisfying at least $(1-\epsilon)m$ constraints.*

Remark 1.2. *The reader may think of the output "typical" as a* certification *that the* 3XOR *instance has no $(1-\epsilon)$-satisfying solution. Note that with high probability the random* 3XOR *instance will not even have a .51-satisfying solution. Feige originally stated his hypothesis for the random 3SAT problem rather than the random 3XOR problem, but he showed the conjectures are equivalent. See also the work of Alekhnovich [Ale03].*

Feige's R3XOR Hypothesis is a fairly well-believed conjecture. The variation in which the XOR constraints may involve any number of variables (not just 3) is called LPN (Learning Parities with Noise) and is believed to be hard even with any $m = \text{poly}(n)$ constraints. The further variation which has linear equations modulo a large prime rather than modulo 2 is called LWE (Learning With Errors) and forms the basis for a large body of cryptography. (See [Reg05] for more on LPN and LWE.) The fastest known algorithm for solving Feige's R3XOR problem seems to be the $2^{O(n/\log n)}$-time algorithm of Blum, Kalai, and Wasserman [BKW03]. Thus it's plausible that RobustGiso requires $2^{n^{1-o(1)}}$ time, which would make it a much more difficulty problem than Giso.

In fact, we are able to base Theorem 1.2 on a larger class of hypotheses which generalize Feige's R3XOR Hypothesis. This is inspired by the recent work of Chan [Cha13], who showed approximation resistance for all predicates supported on a pairwise independent abelian subgroup. In light of this, we show that if no polynomial-time algorithm can distinguish random instances from nearly satisfiable instances of *any* such predicate, then Theorem 1.2 holds. Proving Theorem 1.2 in this more general setting requires some additional tricks. For example, we give a novel gadget construction whose only automorphisms correspond to group elements. See Section 7 for more details.

We close by mentioning some related literature on approximate graph isomorphism. The problem of finding a vertex permutation which maximizes the number of edge overlaps (or minimizes the number of edge/nonedge overlaps) was perhaps first discussed by Arora, Frieze, and Kaplan [AFK02]. They gave an additive PTAS in the case of dense graphs ($m = \Omega(n^2)$). Arvind et al. [AKKV12] recently defined and studied several variants of the approximate graph isomorphism problem. Some of their results concern the case in which G and H have noticeably different numbers of edges and one isn't "punished" for uncovered edges in H. This kind of variant is more like approximate *subgraph* isomorphism, and is much harder. (E.g., when G is a k-clique and H is a general graph the problem is roughly equivalent to the notorious Densest-kSubgraph problem.) The result of theirs which is most relevant to the present work involves hardness of finding approximate isomorphisms in *colored* graphs. In particular, Arvind eta al. prove the following:

Theorem 1.3. *([AKKV12].) There is a linear-time reduction from 2XOR (modulo 2) instances $\mathcal{I}$ to pairs of* colored *graphs G, H such that G and H are α-isomorphic if and only if $\mathcal{I}$ has a solution satisfying at least an α-fraction of constraints. In particular it is* NP*-hard to approximate α-isomorphism for colored graphs to a factor exceeding $\frac{11}{12}$ and* UniqueGames*-hard to approximate it to a factor exceeding .878 (by results of [TSSW00, Hås01], [KKMO07, MOO10] respectively).*

In particular, the theorem holds for colored graphs in which each color class contains at most 4 vertices. However, we do not see any way of eliminating the colors and getting the analogous inapproximability results for the usual Giso problem without using gadgets that would destroy the constant-factor gap.

1.3 Robust asymmetry of random graphs

One of our main technical contributions is showing that random graphs are "robustly asymmetric". In doing so, we generalize the concept of an *asymmetric graph*, which is a graph whose only automorphism is the trivial identity automorphism. A line of research (see, e.g., [ER63, Bol82, MW84, KSV02]) has shown that several distributions of random graphs produce asymmetric graphs with high probability. In their well-known $\mathcal{G}(n,p)$ model, Erdős and Rényi [ER63] proved that for $\frac{\ln n}{n} \leq p \leq 1 - \frac{\ln n}{n}$, $\mathcal{G}(n,p)$ is asymmetric with high probability. If we instead consider a uniformly random d-regular n-vertex graph, the sequence of works [Bol82, MW84, KSV02] shows that we get an asymmetric graph with high probability for any $3 \leq d \leq n-4$. In this work we will work with a third variant, the $\mathcal{G}_{n,m}$ model, in which a graph is chosen uniformly at random from all simple graphs with n vertices and m edges.

Given a graph G and a permutation π over $V(G)$, we call π an α-automorphism if the application of π on G preserves at least an α fraction of the edges. A graph G is (β,γ)-asymmetric if any γ-automorphism π has more than a fraction of $(1-\beta)$ fixed points. Intuitively, when $\beta = 1/n$, $\gamma = 0$, the property is exactly the asymmetry property; when β and γ become larger, the property requires that any permutation that is far from identity is far from an automorphism for the graph. We encourage the reader to refer to Section 2 for the precise definitions.

In this paper, we show the following robust asymmetric property of $\mathcal{G}_{n,m}$.

Theorem 1.4. *For large enough n, suppose that $m = cn$, where $10^4 \leq c \leq n/10^{10}$. Let $\beta^* = \max\{e^{-c/6}, \frac{1}{n}\}$. With probability $(1-n^{-15})$, for all β such that $\beta^* \leq \beta \leq 1$, $\mathcal{G}_{n,m}$ is $(\beta, \beta/240)$-asymmetric.*

A couple of comments are in order. First, an $\exp(-O(c))$ lower bound on β is necessary. This is because there are at least $\lfloor \exp(-O(c)) \cdot n \rfloor$ isolated vertices in $\mathcal{G}_{n,m}$ with high probability. The permutations which only permute these isolated vertices are 1-automorphisms. Therefore, with high probability, $\mathcal{G}_{n,m}$ is not $(\exp(-\omega(c)), 0)$-asymmetric. Second, it is possible to extend our theorem to the $\mathcal{G}(n,p)$ model by showing that there exists a constant $C > 0$, such that for $\frac{C}{n} < p < \frac{1}{C}$, with high probability, $\mathcal{G}(n,p)$ is $(\beta, \beta/240)$-asymmetric for all $\beta \geq \max\{\exp(-\frac{pn}{C}), \frac{1}{n}\}$. Third, when $c \geq 6 \ln n$ (or, when $p \geq \frac{C \ln n}{n}$ in the $\mathcal{G}(n,p)$ model), we can let $\beta = \frac{1}{n}$ and obtain that $\mathcal{G}_{n,m}$ ($\mathcal{G}(n,p)$, respectively) is asymmetric with high probability – a result in the flavor of [ER63]. Finally, we do not work hard to optimize the constants in the theorem statement; we believe a more careful analysis would bring them down to more civilized numbers, but it is still interesting to explore the limits of these constants.

Now we briefly explain our proof techniques. Let us consider the case where c is a big constant and $\beta = 1$, so that we only need to worry about the permutations without fixed points. We would like to show that, for every such permutation π, $\Pr_{G \sim \mathcal{G}_{n,m}}[\pi \text{ is a } \frac{1}{240}\text{-automorphism for } G] \ll \frac{1}{n!}$, and therefore we can union bound over all such possible permutations. In order to do this, from all $\binom{n}{2}$ possible edges, we construct $\Omega(n^2)$ disjoint pairs of edges (e, e'), which we will refer to as "bins", such that $\pi(e) = e'$. We call a bin "half-full" if exactly one edge in the pair is selected in G. It is easy to see that whenever there are more than $\frac{m}{120}$ half-full bins, π cannot be a $\frac{1}{240}$-automorphism. At this point, we would like for $\Pr_G\left[\#\text{half-full bins} < \frac{m}{120}\right] \ll \frac{1}{n!}$, and this is easy to show. Unfortunately, this method does not work when $\beta = \frac{1}{2}$. To see why, let π be a permutation with $\frac{n}{2}$ fixed points. The probability that every edge in G has fixed points of π for its endpoints is roughly $2^{-2m} = 2^{-O(n)}$. Therefore we have $\Pr_G[\pi \text{ is an automorphism for } G] \geq 2^{-O(n)}$, and this is not enough for the application of union bound (since there are more than $(n/2)! = 2^{\Omega(n \log n)}$ such permutations). A possible fix to this problem is: we first show that with high probability $(1-n^{-\omega(1)})$, for every π with $\frac{n}{2}$ fixed points, there are many edges of G with at least one end point not fixed by π; then, conditioned on this event, we show the probability that a fixed π is not a $\frac{1}{480}$-automorphism is small enough for the union bound method. The actual proof is more involved, and it is also technically challenging to work with c as large as $\Omega(n)$, and β as small as $\frac{1}{n}$.

Finally, for our application to the GRAPHISOMORPHISM problem, we need to extend Theorem 1.4 to hypergraphs. More details on robust asymmetry of random hypergraphs can be found in Section 6.

1.4 Organization of the paper

In Section 2, we introduce the notations and the SOS/Lasserre hierarchy. In Section 3, we describe a reduction from 3XOR to GISO. The completeness and soundness lemmas for reduction are proved in Section 4 and Section 5 respectively. In Section 6, we prove robust asymmetry property for random graphs and random hyper-

graphs. In Section 7, we generalize our RobustGIso hardness result to be based on a larger class of hypotheses. We discuss some of the future directions in Section 8.

Proofs of the main theorems. Theorem 1.1 follows from Theorem 2.1, Lemma 4.2 and Lemma 5.1, by choosing $c = 10^6$. Theorem 1.2 follows from Lemma 4.1 and Lemma 5.1, by choosing $c = \max\{10^6, \Delta\}$.

2 Preliminaries

We will be working with undirected graphs and hypergraphs, both of which will be denoted by $G = (V, E)$. Here, an undirected edge $e \in E$ is a set of 2 vertices $\{i, j\}$ for graphs and a set of k vertices $\{i_1, i_2, \ldots, i_k\}$ for k-uniform hypergraphs. When G is an directed graph, we use (i, j) to denote a directed edge. We also use the notation $V(G)$ to denote the vertex set of G, and $E(G)$ to denote the edge set of G.

For any two undirected graphs (or hypergraphs) $G = (V(G), E(G))$ and $H = (V(H), E(H))$ with the same number of vertices, and for any bijection $\pi : V(G) \to V(H)$, let

$$\mathsf{GI}(G, H; \pi) = \frac{|\{e \in E(G) : \pi(e) \in E(H)\}|}{\max\{|E(G), E(H)|\}},$$

where $\pi(e)$ is the edge obtained by applying π on each vertex incident to e. Let

$$\mathsf{GI}(G, H) = \max_{\pi : V(G) \to V(H)} \mathsf{GI}(G, H; \pi).$$

We say that an edge $e \in E(G)$ is *satisfied* by π if $\pi(e) \in E(H)$. We call π an α-*isomorphism* for G and H if $\mathsf{GI}(G, H; \pi) \geq \alpha$, and we say G and H are α-*isomorphic* if $\mathsf{GI}(G, H) \geq \alpha$.

For any permutation $\pi : V(G) \to V(G)$, let

$$\mathsf{AUT}(G; \pi) = \mathsf{GI}(G, G; \pi).$$

We say that π is an α-*automorphism* for G if $\mathsf{AUT}(G; \pi) \geq \alpha$.

Given a permutation π over the set V, an element $i \in V$ is a *fixed point* of π if $\pi(i) = i$.

Definition 3. *A graph (possibly hypergraph) G is* (β, γ)-asymmetric *if, for any permutation π on the vertex set of G that has at most $(1 - \beta)$ fraction of the vertices as fixed points, we have $\mathsf{AUT}(G; \pi) < 1 - \gamma$.*

We extend the $\mathcal{G}_{n,m}$ random graph model to hypergraphs as follows. Let $\mathcal{G}^{(k)}_{n,m}$ be the uniform distribution over all $\binom{\binom{n}{k}}{m}$ simple k-uniform hypergraphs with n vertices and m edges.

A 3XOR instance $\mathcal{C}$ is a collection of equations $C_1, C_2, \ldots, C_m$ over the variable set $\mathcal{X}$. Each equation C_i is of the form $x_{j_1} + x_{j_2} + x_{j_3} = b$ where $x_{j_1}, x_{j_2}, x_{j_3}$ are the variables from $\mathcal{X}$, $b \in \mathbb{Z}_2$. Given an assignment $\tau : \mathcal{X} \to \mathbb{Z}_2$, let $\mathsf{val}(\mathcal{C}; \tau)$ be the fraction of equations in $\mathcal{C}$ satisfied by τ. Let $\mathsf{val}(\mathcal{C}) = \max_{\tau : \mathcal{X} \to \mathbb{Z}_2} \mathsf{val}(\mathcal{C}; \tau)$.

2.1 SOS/Lasserre hierarchy

One way to formulate the SOS/Lasserre hierarchy is via the pseudo-expectation view. We briefly recall the formulation as follows. More discussion about this view can be found in [BBH+12].

We consider the feasibility of a system over n variables $(x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$ with the following constraints: $P_i(x) = 0$ for $i = 1, 2, \ldots, m_P$ and $Q_i(x) \geq 0$ for $j = 1, 2, \ldots, m_Q$, where all the P_i, Q_i polynomials are of degree at most d. For $r \geq d$, the degree-r SOS/Lasserre hierarchy finds a pseudo-expectation operator $\widetilde{\mathbf{E}}[\cdot]$ defined on the space of real polynomials of degree at most r over intermediates $x_1, x_2, \ldots, x_n$ such that

- $\widetilde{\mathbf{E}}[1] = 1$;
- $\widetilde{\mathbf{E}}[\alpha p + \beta q] = \alpha \widetilde{\mathbf{E}}[p] + \beta \widetilde{\mathbf{E}}[q]$ for all real numbers α, β and all polynomials p and q of degree at most r;
- $\widetilde{\mathbf{E}}[p^2] \geq 0$ for all polynomials p of degree at most $r/2$;
- $\widetilde{\mathbf{E}}[P_i \cdot q] = 0$ for all $i = 1, 2, \ldots, m_P$ and all polynomials q such that $P_i \cdot q$ is of degree at most r;
- $\widetilde{\mathbf{E}}[Q_i \cdot p^2] \geq 0$ for all $i = 1, 2, \ldots, m_Q$ and all polynomials p such that $Q_i \cdot p^2$ is of degree at most r.

We call any operator $\widetilde{\mathbf{E}}[\cdot]$ a normalized linear operator if it has the first two properties listed above.

SOS/Lasserre hierarchy for 3XOR**.** Let $\mathcal{C}$ be a 3XOR instance on variable set $\mathcal{X}$. The degree-r SOS/Lasserre hierarchy for the natural integer programming for (the satisfiability of) $\mathcal{C}$ is to find a normalized linear pseudo-expectation operator $\widetilde{\mathbf{E}}[\cdot]$ defined on the space of polynomials of degree at most r over the indeterminates $(A[x \mapsto a])_{x\in\mathcal{X},a\in\mathbb{Z}_2}$ associated to $\mathcal{C}$, such that

1. $\widetilde{\mathbf{E}}[(A[x \mapsto a]^2 - A[x \mapsto a]) \cdot q] = 0$ for all $x \in \mathcal{X}$, $a \in \mathbb{Z}_2$, and polynomials q;
2. $\widetilde{\mathbf{E}}[(A[x \mapsto 0] + A[x \mapsto 1] - 1) \cdot q] = 0$ for all $x \in \mathcal{X}$ and polynomials q;
3. $\widetilde{\mathbf{E}}[(\sum_{\alpha_C \text{ satisfying } C} A[x_1 \mapsto \alpha_C(x_1)]A[x_2 \mapsto \alpha_C(x_2)]A[x_3 \mapsto \alpha_C(x_3)] - 1) \cdot q] = 0$ for each $C \in \mathcal{C}$ involving variables x_1, x_2, x_3 and all polynomials q;
4. $\widetilde{\mathbf{E}}[p^2] \geq 0$ for all polynomials p.

We say there is degree-r SOS refutation for the satisfiability of $\mathcal{C}$ if the pseudo-expectation operator with properties listed above does not exist.

SOS/Lasserre hierarchy for GISO**.** Let $G = (V(G), E(G))$ and $H = (V(H), E(H))$ be two graphs such that $|V(G)| = |V(E)|$, $|E(G)| = |E(H)|$. The degree-r SOS/Lasserre hierarchy for the natural integer programming the isomorphism problem between G and H is to find a normalized linear pseudo-expectation operator $\widetilde{\mathbf{E}}[\cdot]$ on the space of real polynomials of degree at most r over the indeterminates $(\Pi[u \mapsto v])_{u\in V(G),v\in V(H)}$ such that:

a. $\widetilde{\mathbf{E}}[(\Pi[u \mapsto v]^2 - \Pi[u \mapsto v]) \cdot q] = 0$ for all $u \in V(G)$, $v \in V(H)$, and polynomials q;

b. $\widetilde{\mathbf{E}}[(\sum_{v\in V(H)} \Pi[u \mapsto v]-1)\cdot q] = 0$ for all $u \in V(G)$ and polynomials q; and similarly, $\widetilde{\mathbf{E}}[(\sum_{u\in V(G)} \Pi[u \mapsto v] - 1) \cdot q] = 0$ for all $v \in V(H)$ and polynomials q;

c. $\widetilde{\mathbf{E}}[(\sum_{\{u,u'\}\in E(G)} \sum_{v,v':\{v,v'\}\in E(H)} \Pi[u \mapsto v]\Pi[u' \mapsto v'] - |E(G)|) \cdot p^2] \geq 0$ for all polynomials p;

d. $\widetilde{\mathbf{E}}[p^2] \geq 0$ for all polynomials p.

We say there is degree-r SOS refutation for the isomorphism between G and H if the pseudo-expectation operator with properties listed above does not exist.

Remark 2.1. *It is equivalent to replace (c) by "*$\widetilde{\mathbf{E}}[(\sum_{v,v':\{v,v'\}\in E(H)} \Pi[u \mapsto v]\Pi[u' \mapsto v'] - 1) \cdot q] = 0$ *for all* $(u, u') \in E(G)$ *and all polynomials* q*".*

2.2 Random 3XOR

A random 3XOR instance with n variables and m equations is sampled by choosing m unordered 3-tuples of variables from all possible $\binom{n}{3}$ ones, and making each 3-tuple $(x_{j_1}, x_{j_2}, x_{j_3})$ into a 3XOR constraint $x_{j_1} + x_{j_2} + x_{j_3} = b$ with an independent random $b \in \mathbb{Z}_2$.

Theorem 2.1. *[Sch08] For every constant $c > 1$, there is exists $\eta > 0$ such that with probability $1 - o(1)$, the satisfiability of a random* 3XOR *instance*[4] *with n variables and cn equations cannot be refuted by degree-(ηn) SOS/Lasserre hierarchy.*

3 Reduction from 3XOR to GISO

We define a slight variant of the basic gadget from [CFI92]:

Definition 4. *Let C be a* 3XOR *constraint involving variables x_1, x_2, x_3. The associated gadget graph G_C consists of: 6 "variable vertices" with names "$x_i \mapsto a$" for each $i \in [3]$, $a \in \mathbb{Z}_2$; and, 4 "constraint vertices" with names "$x_1 \mapsto a_1, x_2 \mapsto a_2, x_3 \mapsto a_3$" for each partial assignment which satisfies the constraint C. Regarding edges, each pair of variable vertices $x_i \mapsto 0$, $x_i \mapsto 1$ is connected by an edge; the four constraint vertices are connected by a clique; and, each constraint vertex α is connected to the three variable vertices it is consistent with.*

[4]The random 3XOR distribution used in [Sch08] is slightly different, but the theorem still holds for our distribution.

Now we describe how an entire instance of 3XOR is encoded by a graph:

Definition 5. *Let $\mathcal{C}$ be a collection of* 3XOR *constraints over variable set $\mathcal{X}$. We define the associated graph $G_{\mathcal{C}}$ as follows: For each constraint $C \in \mathcal{C}$, the graph $G_{\mathcal{C}}$ contains a copy of the gadget graph G_C. However we* identify *all of the variable vertices $x \mapsto a$ across $x \in \mathcal{X}$, $a \in Z_2$ as well as the variable edges $(x \mapsto 0, x \mapsto 1)$. The constraint vertices associated to C, on the other hand, are left as-is, and will be named α_C. We denote the set of vertices $\{x \mapsto 0, x \mapsto 1\}$ by V_x for every variable x, denote the set of vertices corresponding to C by V_C for every variable C.*

Remark 3.1. *If $\mathcal{C}$ is a* 3XOR *instance with n vertices and m constraints then the graph $G_{\mathcal{C}}$ has $N = 4m + 2n$ vertices and $M = 18m + n$ edges.*

Finally, we introduce the following notation:

Notation 3.1. *Let C be a* 3XOR *constraint involving variables x_i, x_j, x_k. We write $\underline{C}$ for its homogeneous version, $x_i + x_j + x_k = 0$. Given a collection of* 3XOR *constraints $\mathcal{C}$ we write $\underline{\mathcal{C}} = \{\underline{C} : C \in \mathcal{C}\}$.*

The reduction. *Given a collection of* 3XOR *constraints $\mathcal{C}$, the corresponding* GISO *instance is $(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}})$.*

4 Completeness

Lemma 4.1 (Completeness)**.** *If $\mathcal{C}$ is a* 3XOR *instance such that $\mathsf{val}(\mathcal{C}) \geq 1 - \epsilon$, then $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq 1 - 2\epsilon/3$.*

Proof. Let τ be an assignment to the variables in $\mathcal{C}$ such that $\mathsf{val}(\mathcal{I}; \tau) \geq 1 - \epsilon$. Now we define a bijection π from the vertices in $G_{\mathcal{C}}$ to the ones in $G_{\underline{\mathcal{C}}}$ as follows.

For each variable vertex $x_j \mapsto b$, let $\pi(x_j \mapsto b) = x_j \mapsto b + \tau(x_j)$. For any equation vertex α_{C_i}, if C_i is not satisfied by τ, map it to an arbitrary vertex in $V_{\underline{C_i}}$. If C_i is satisfied by τ, let us suppose that $C_i : x_{j_1} + x_{j_2} + x_{j_3} = b$, let α' be an assignment such that $\alpha'(x_{j_t}) = \alpha(x_{j_t}) + \tau(x_{j_t})$ for all $t \in \{1, 2, 3\}$. Observe that

$$\alpha'(x_{j_1}) + \alpha'(x_{j_2}) + \alpha'(x_{j_3}) = (\alpha(x_{j_1}) + \alpha(x_{j_2}) + \alpha(x_{j_3})) + (\tau(x_{j_1}) + \tau(x_{j_2}) + \tau(x_{j_3})) = b + b = 0.$$

Therefore $\alpha'_{\underline{C_i}}$ is a vertex in $G_{\underline{\mathcal{C}}}$. We let π map α_{C_i} to $\alpha'_{\underline{C_i}}$.

It is straightforward to check that all the edges between equation vertices and between variable vertices are satisfied. Now we consider an edge between a equation vertex and a variable vertex, namely between α_{C_i} and $x_j \mapsto b$ where x_j is an variable in equation C_i and $\alpha(x_j) = b$. We show that the edge is satisfied by π whenever C_i is satisfied by τ. Let α' and b' be such that $\pi(\alpha_{C_i}) = \alpha'_{\underline{C_i}}$, $\pi(x_j \mapsto b) = x_j \mapsto b'$. Observe that

$$\alpha'(x_j) = \alpha(x_j) + \tau(x_j) = b + \tau(x_j) = b',$$

and this implies that there is an edge between α'_{C_i} and $x_j \mapsto b'$.

We see that the only edges in $G_{\mathcal{C}}$ which might not be satisfied by π are the ones between equation vertices and variable vertices where the corresponding equation vertex is not satisfied by τ. For each equation not satisfied, there are at most 12 such edges. Therefore there are at most $12\epsilon m$ edges not satisfied. We have

$$\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq \mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}; \pi) \geq \frac{M - 12\epsilon m}{M} \geq 1 - \frac{2}{3}\epsilon. \qquad \square$$

4.1 SOS completeness

Lemma 4.2 (SOS completeness)**.** *Let $\mathcal{C}$ be a* 3XOR *instance on variable set $\mathcal{X}$ and suppose that every SOS refutation of $\mathcal{C}$ requires degree exceeding r. Then every SOS refutation of the statement "$G_{\mathcal{C}}$ and $G_{\underline{\mathcal{C}}}$ are isomorphic" requires degree exceeding $r/3$.*

Proof. Since $\mathcal{C}$ cannot be refuted in degree r, there is a pseudo-expectation operator $\widetilde{\mathbf{E}}_{\mathcal{C}}[\cdot]$ defined on the space of real polynomials of degree at most r over the indeterminates $(A[x \mapsto a])_{x \in \mathcal{X}, a \in \mathbb{Z}_2}$ associated to $\mathcal{C}$. This $\widetilde{\mathbf{E}}_{\mathcal{C}}[\cdot]$ is normalized, linear, and satisfies:

i. $\widetilde{\mathbf{E}}_{\mathcal{C}}[(A[x \mapsto a]^2 - A[x \mapsto a]) \cdot q] = 0$ for all $x \in \mathcal{X}$, $a \in \mathbb{Z}_2$, and polynomials q;

ii. $\widetilde{\mathbf{E}}_{\mathcal{C}}[(A[x \mapsto 0] + A[x \mapsto 1] - 1) \cdot q] = 0$ for all $x \in \mathcal{X}$ and polynomials q;

iii. $\widetilde{\mathbf{E}}_{\mathcal{C}}[(\sum_{\alpha_C \text{ satisfying } C} A[x_1 \mapsto \alpha_C(x_1)]A[x_2 \mapsto \alpha_C(x_2)]A[x_3 \mapsto \alpha_C(x_3)] - 1) \cdot q] = 0$ for each $C \in \mathcal{C}$ involving variables x_1, x_2, x_3 and all polynomials q;

iv. $\widetilde{\mathbf{E}}_{\mathcal{C}}[p^2] \geq 0$ for all polynomials p.

Our task is to define a normalized linear pseudo-expectation operator $\widetilde{\mathbf{E}}_{\mathcal{G}}[\cdot]$ on the space of real polynomials of degree at most $r/3$ over the indeterminates $(\Pi[u \mapsto v])_{u \in V(G_{\mathcal{C}}), v \in V(G_{\underline{\mathcal{C}}})}$ such that:

I. $\widetilde{\mathbf{E}}_{\mathcal{G}}[(\Pi[u \mapsto v]^2 - \Pi[u \mapsto v]) \cdot q] = 0$ for all $u \in V(G_{\mathcal{C}})$, $v \in V(G_{\underline{\mathcal{C}}})$, and polynomials q;

II. $\widetilde{\mathbf{E}}_{\mathcal{G}}[(\sum_{v \in V(G_{\underline{\mathcal{C}}})} \Pi[u \mapsto v] - 1) \cdot q] = 0$ for all $u \in V(G_{\mathcal{C}})$ and polynomials q; and similarly, $\widetilde{\mathbf{E}}_{\mathcal{G}}[(\sum_{u \in V(G_{\mathcal{C}})} \Pi[u \mapsto v] - 1) \cdot q] = 0$ for all $v \in V(G_{\underline{\mathcal{C}}})$ and polynomials q;

III. $\widetilde{\mathbf{E}}_{\mathcal{G}}[(\sum_{\{u,u'\} \in E(G_{\mathcal{C}})} \sum_{v,v' : \{v,v'\} \in E(G_{\underline{\mathcal{C}}})} \Pi[u \mapsto v]\Pi[u' \mapsto v'] - M) \cdot p^2] \geq 0$ for all polynomials p;

IV. $\widetilde{\mathbf{E}}_{\mathcal{G}}[p^2] \geq 0$ for all polynomials p.

Here M denotes the number of edges in $G_{\mathcal{C}}$ (and also in $G_{\underline{\mathcal{C}}}$).

The idea is to formally define each indeterminate $\Pi[u \mapsto v]$ as a certain degree-3 multilinear polynomial of the indeterminates $A[x \mapsto a]$. Then $\widetilde{\mathbf{E}}_{\mathcal{G}}[\cdot]$ is automatically defined in terms of $\widetilde{\mathbf{E}}_{\mathcal{C}}[\cdot]$ for all polynomials of degree at most $r/3$. The natural definition is as follows:

1. Let $x \in \mathcal{X}$ and $a \in \mathbb{Z}_2$. We define $\Pi[(x \mapsto a) \mapsto (x \mapsto b)] = A[x \mapsto (a - b)]$.
2. Let $C \in \mathcal{C}$, let $\alpha_C = (x_1 \mapsto a_1, x_2 \mapsto a_2, x_3 \mapsto a_3)$ be constraint vertex in $G_{\mathcal{C}}$ corresponding to C, and let $\beta_{\underline{C}} = (x_1 \mapsto b_1, x_2 \mapsto b_2, x_3 \mapsto b_3)$ be a constraint vertex in $G_{\underline{\mathcal{C}}}$ corresponding to $\underline{C}$. We define $\Pi[\alpha_C \mapsto \beta_{\underline{C}}]$ to be the following degree-3 monomial:
$$A[x_1 \mapsto (a_1 - b_1)]A[x_2 \mapsto (a_2 - b_2)]A[x_3 \mapsto (a_3 - b_3)].$$
3. All other indeterminates $\Pi[u \mapsto v]$ are formally defined to be 0.

It is clear that $\widetilde{\mathbf{E}}_{\mathcal{G}}[\cdot]$ is normalized and linear by the same property of $\widetilde{\mathbf{E}}_{\mathcal{C}}[\cdot]$. It remains to show that the induced pseudo-expectation operator $\widetilde{\mathbf{E}}_{\mathcal{G}}[\cdot]$ satisfies (I)–(IV) using the fact that $\widetilde{\mathbf{E}}_{\mathcal{C}}[\cdot]$ satisfies (i)–(iv). Most of these are easy; for example, the implication (iv) $\Rightarrow$ (IV) is immediate. Almost as easy is that (i) $\Rightarrow$ (I) and that (i), (ii) $\Rightarrow$ (II). We illustrate some of these implications, leaving the rest to the reader. For example, let's verify (I) for indeterminates of type $\Pi[\alpha_C \mapsto \beta_{\underline{C}}]$. For brevity we'll write $\Pi[\alpha_C \mapsto \beta_{\underline{C}}]$ as $A_1A_2A_3$. Now for any polynomial q over the Π's,

$$\begin{aligned}
&\widetilde{\mathbf{E}}_{\mathcal{G}}[(\Pi[\alpha_C \mapsto \beta_{\underline{C}}]^2 - \Pi[\alpha_C \mapsto \beta_{\underline{C}}]) \cdot q] \\
&= \widetilde{\mathbf{E}}_{\mathcal{G}}[(A_1^2A_2^2A_3^2 - A_1A_2A_3) \cdot q'] && \text{(for some polynomial } q' \text{ over the } A\text{'s)} \\
&= \widetilde{\mathbf{E}}_{\mathcal{G}}[(A_1^2 - A_1)A_2^2A_3^2 \cdot q'] + \widetilde{\mathbf{E}}_{\mathcal{G}}[A_1(A_2^2 - A_2)A_3^2 \cdot q'] + \widetilde{\mathbf{E}}_{\mathcal{G}}[A_1A_2(A_3^2 - A_3) \cdot q'] \\
&= 0 && \text{(by (i).)}
\end{aligned}$$

And let's verify (II) when u is a variable vertex $x \mapsto a$:

$$\begin{aligned}
&\widetilde{\mathbf{E}}_{\mathcal{C}}\left[\left(\sum_{v \in V(G_{\underline{\mathcal{C}}})} \Pi[(x \mapsto a) \mapsto v] - 1\right) \cdot q\right] \\
&= \widetilde{\mathbf{E}}_{\mathcal{C}}[(A[x \mapsto a - 0] + A[x \mapsto a - 1] - 1) \cdot q'] && \text{(all other } \Pi\text{'s are 0)} \\
&= 0. && \text{(by (ii))}
\end{aligned}$$

The main effort is to establish (III). In fact we will show

$$\widetilde{\mathbf{E}}_{\mathcal{C}}\left[\left(\sum_{v,v' : \{v,v'\} \in E(G_{\underline{\mathcal{C}}})} \Pi[u \mapsto v]\Pi[u' \mapsto v'] - 1\right) \cdot p^2\right] = 0 \tag{1}$$

for all edges $\{u, u'\} \in E(G_{\mathcal{C}})$ and all p, whence (III) follows by summing. We will omit the (easy) verification of this for the edges $(x \mapsto 0, x \mapsto 1)$. Instead we will first verify that (1) holds for a typical clique edge associated to constraint C, say (α_C, α'_C), on variables x_1, x_2, x_3. Only the indeterminates of corresponding constraints, say $\Pi[\alpha_C \mapsto \beta_{\underline{C}}]$, are nonzero. Writing $A_i[\alpha - \beta] = A[x_i \mapsto \alpha_C(x_i) - \beta_{\underline{C}}(x_i)]$ for brevity (and similarly with primes), the quantity in (1) is

$$\widetilde{\mathbf{E}}_{\mathcal{C}}\Big[\Big(\sum_{\substack{\beta_{\underline{C}},\beta'_{\underline{C}}\\ \text{satisfying } \underline{C}}} A_1[\alpha-\beta]A_2[\alpha-\beta]A_3[\alpha-\beta]A_1[\alpha'-\beta']A_2[\alpha'-\beta']A_3[\alpha'-\beta'] - 1\Big)\cdot p'^2\Big]$$

$$= \widetilde{\mathbf{E}}_{\mathcal{C}}\Big[\Big(\big(\sum_{\beta_{\underline{C}}} A_1[\alpha-\beta]A_2[\alpha-\beta]A_3[\alpha-\beta]\big)\big(\sum_{\beta'_{\underline{C}}} A_1[\alpha'-\beta']A_2[\alpha'-\beta']A_3[\alpha'-\beta']\big) - 1\Big)\cdot p'^2\Big]. \quad (2)$$

Now for fixed α_C, as $\beta_{\underline{C}}$ ranges over all satisfying assignments to $\underline{C}$, the assignment $\alpha_C - \beta_{\underline{C}}$ ranges over all satisfying assignments to C. The analogous statement holds also for α'_C. It's now straightforward to see that the vanishing of (2) follows from (iii).

Our final task is to verify (1) also for edges between variable vertices and constraint vertices. Fix a typical such edge, say one connecting $x_1 \mapsto a_1$ to α_C. (We'll use the same notation as before for α_C; in particular, note that we must have $\alpha_C(x_1) = a_1$.) Now in this case, the quantity in (1) is

$$\widetilde{\mathbf{E}}_{\mathcal{C}}\Big[\Big(\sum_{\substack{b\in\mathbb{Z}_2,\\ \beta_{\underline{C}} \text{ satisfying } \underline{C}}} A[x_1 \mapsto (a_1 - b)]A_1[\alpha-\beta]A_2[\alpha-\beta]A_3[\alpha-\beta] - 1\Big)\cdot p'^2\Big]$$

$$= \widetilde{\mathbf{E}}_{\mathcal{C}}\Big[\Big(\big(\sum_{c\in\mathbb{Z}_2} A[x_1 \mapsto c]\big)\big(\sum_{\beta_{\underline{C}}} A_1[\alpha-\beta]A_2[\alpha-\beta]A_3[\alpha-\beta]\big) - 1\Big)\cdot p'^2\Big]. \quad (3)$$

Again, the fact that (3) vanishes now easily follows from (ii), (iii). □

5 Soundness

In this section, we prove the following soundness lemma.

Lemma 5.1 (Soundness)**.** *Let $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$ be a random* 3XOR *instance with n variables and $m = cn$ equations where $c \geq 10^10$. With probability $1 - o(1)$, we have*

$$\mathsf{GI}(G_C, G_{\underline{C}}) < 1 - \frac{1}{95c^2}.$$

Before proving Lemma 5.1, we first introduce the following definition.

Definition 6. *A graph (possibly hypergraph) G is* (ϵ, D)-degree bounded *if the average degree of every set of ϵ fraction of vertices is at most D.*

Claim 5.2. *Suppose $c \geq 3$. A random 3-uniform hypergraph H drawn from $\mathcal{G}^{(3)}_{n,m}$, where $m = cn$, is $(1/c, 100c)$-degree bounded with probability $1 - o(1)$.*

Claim 5.3. *Given an (ϵ, D)-degree bounded graph G with n vertices, every set of βn vertices has at most $(\epsilon + \beta)Dn$ edges incident to them.*

Lemma 5.1 is directed implied by the following two lemmas.

Lemma 5.4. *Let $H = ([n], E = \{e_i\})$ be a 3-uniform hypergraph with n vertices and $m = cn$ hyperedges. the constraint graph of a* 3XOR *instance with n variables and $m = cn$ equations. Suppose H is $(\epsilon, 100c)$-degree bounded and (β, γ)-asymmetric, where $\gamma \geq 200\epsilon$. Let $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$ be an arbitrary* 3XOR *with n variables and m constraints based on H. (In other words, each hyperedge e_i of H connects the indices of the* 3 *variables used by C_i.) If we set*

$$\delta := \delta(c, \epsilon, \beta, \gamma) = \min\left\{\frac{1}{200}, \frac{\gamma}{48}, \frac{\epsilon}{95c}\right\},$$

when $\mathsf{GI}(G_C, G_{\underline{C}}) \geq 1 - \delta$, we have $\mathsf{val}(\mathcal{C}) \geq .9 - 100(\epsilon + \beta)$.

Lemma 5.5. *If $\mathcal{C}$ is a random* 3XOR *instance with n variables and $m \geq 10000n$ equations, then with probability $1 - o(1)$, we have $\mathsf{val}(\mathcal{C}) < .51$.*

Proof of Lemma 5.1. Set $\epsilon = \frac{1}{c}$, $\gamma = \frac{200}{c}$, $\beta = \frac{48000}{c}$. Combining Claim 5.2, Lemma 5.5, and Theorem 6.9, we know that with probability $(1 - o(1))$, all of the following hold:

1. H is $(\epsilon, 100c)$-degree bounded,
2. H is (β, γ)-asymmetric, and
3. $\mathsf{val}(\mathcal{C}) < .51$.

Given that these hold, assume for sake of contradiction that $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq 1 - \frac{1}{95c^2}$. Then because G satisfies Item 1 and Item 2, Lemma 5.4 implies that

$$\mathsf{val}(\mathcal{C}) \geq .9 - 100\left(\frac{1}{c} + \frac{48000}{c}\right) \geq .8,$$

where the last step follows because $c \geq 10^{10}$. However, this contradicts Item 3. Therefore, $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) < 1 - \frac{1}{95c^2}$ with probability $1 - o(1)$. □

The proof of Lemma 5.5 is standard.

Proof of Lemma 5.5. Fix an assignment to the n variables. The probability that the assignment satisfies at least $.51m$ equations of a random 3XOR instance is at most $\exp(-.02^2 \cdot .5m/2) = \exp(-.0001m)$ by the Chernoff bound. Since there are only 2^n assignments, the probability that no assignment satisfies more than $.5m$ equations is at least $1 - \exp(-.0001m) \cdot 2^n = 1 - o(1)$ when $m \geq 10000n$. □

The rest of the section is devoted to the proof of Lemma 5.4.

Proof of Lemma 5.4. Let π be a bijection mapping the vertices in $G_{\mathcal{C}}$ to the vertices in $G_{\underline{\mathcal{C}}}$ such that $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}; \pi) \geq 1 - \delta$. We first prove that for most i's, π maps the set V_{C_i} to $V_{\underline{C_{i'}}}$ for some i', and for most j's, π maps most V_{x_j} to $V_{x'_j}$ for some j'. Formally, let A be the set of $i \in [m]$ such that $\pi(V_{C_i}) = V_{\underline{C_{i'}}}$ for some i', and let B be the set of $j \in [n]$ such that $\pi(V_{x_j}) = V_{x_{j'}}$ for some j'. We show that

Claim 5.6. $|A| \geq (1 - 19\delta)m$, $|B| \geq (1 - 95c\delta)n$.

Now we are able to define a permutation σ on the variables in $\mathcal{C}$ (as well as $\underline{\mathcal{C}}$ since the set of variables is shared). We let σ to be an arbitrary permutation so that for each $j \in B$, we have $\sigma(j) = j'$ where $\pi(V_{x_j}) = V_{x_{j'}}$. Now we show that σ is an almost automorphism for the constraint graph (i.e. the hypergraph $H = ([n], E)$).

Claim 5.7. $\mathsf{AUT}(H; \sigma) \geq (1 - 100\epsilon - 24\delta)m$.

By our setting of δ, we have $24\delta < \gamma/2$. Since we also assume that $100\epsilon < \gamma/2$, we have $\mathsf{AUT}(G; \sigma) \geq (1 - \gamma)m$, and therefore we know that σ has at least $(1 - \beta)n$ fixed points.

Now we are ready to define an assignment $\tau : \{x_j\} \to \mathbb{Z}_2$ which certifies that $\mathsf{val}(\mathcal{I}) \geq .9$. For each j which is not a fixed point of σ, define $\tau(x_j)$ arbitrarily. For each j being a fixed point of σ, we know that $\pi(V_{x_j}) = V_{x_j}$. We let $\tau(x_j) = b$ where $\pi(x_j \mapsto 0) = x_j \mapsto b$. We conclude the proof by showing the following claim.

Claim 5.8. $\mathsf{val}(\mathcal{C}; \tau) \geq .9 - 100(\epsilon + \beta)$.

□

5.1 Proof of the claims

Proof of Claim 5.6. Observe that the only 4-cliques in $G_{\underline{\mathcal{C}}}$ are $V_{\underline{C_{i'}}}$ ($i' \in [m]$). Therefore, if $\pi(V_{C_i}) \neq V_{\underline{C_{i'}}}$ for every i', we know that at least one of the edges in the clique V_{C_i} is not satisfied. Therefore we have $m - |A| \leq \delta M$ (recall that $M = 18m + n$ is the number of edges in $G_{\mathcal{C}}$), i.e. $|A| \geq m - \delta M \geq (1 - 19\delta)m$.

We now know that at least $(1 - 19\delta)m \cdot 4$ equation vertices in $G_{\underline{\mathcal{C}}}$ are mapped from equation vertices in $G_{\mathcal{C}}$. Therefore there are at most $4m - (1 - 19\delta)m \cdot 4 = 76\delta m$ equation vertices in $G_{\underline{\mathcal{C}}}$ being mapped from variable vertices in $G_{\mathcal{C}}$. In other words, π maps at least $2n - 76\delta m = (2 - 76c\delta)n$ variable vertices to variable vertices. Let B' be the set of j's such that both vertices in V_{x_j} is mapped to a variable vertex. We have $B' \supseteq B$ and $|B'| \geq (2 - 76c\delta)n - n = (1 - 76c\delta)n$. For each $j \in B' \setminus B$, we know that the edge in V_{x_j} is not satisfied. Therefore $|B' \setminus B| \leq \delta M$. Therefore, $|B| = |B'| - |B' \setminus B| \geq (1 - 76c\delta)n - \delta M \geq (1 - 95c\delta)n$. □

Proof of Claim 5.7. Let E' be the set of hyperedges in E whose vertices are all in B. Since G is $(\epsilon, 100c)$-degree bounded, and by our setting of parameters $95c\delta < \epsilon$, by Claim 5.3, we know that $|E'| \geq m - 100c\epsilon n = (1-100\epsilon)m$. Now let us consider the hyperedges in $E'' = E' \cap A$. (Also observe that $|E''| \geq (1-100\epsilon-19\delta)m$.) We claim that most of the hyperedges in E'' are satisfied by σ. For every hyperedge $e_i = \{j_1, j_2, j_3\} \in E''$ that is not satisfied, we know that $\{\sigma(j_1), \sigma(j_2), \sigma(j_3)\} \notin T$. Since $i \in E'' \subseteq A$, let i' be the equation index such that $\pi(V_{C_i}) = V_{\underline{C_{i'}}}$. Since $\{\sigma(j_1), \sigma(j_2), \sigma(j_3)\} \notin E$, we have $e_{i'} \neq \{\sigma(j_1), \sigma(j_2), \sigma(j_3)\}$. Let us assume w.l.o.g. that $\sigma(j_1) \notin e_{i'}$. Then there is no edge between $V_{\underline{C_{i'}}}$ and $V_{x_{\sigma(j_1)}}$ in $G_{\underline{\mathcal{C}}}$. Therefore the 4 edges between V_{C_i} and $V_{x_{j_1}}$ in $G_{\mathcal{C}}$ are not satisfied.

We have proved that whenever there is an hyperedge in E'' not satisfied by σ, there are at least 4 edges in $G_{\mathcal{C}}$ not satisfied by π. Since π satisfies $(1 - \delta)M$ edges, there are at most $\delta M/4$ hyperedges in E'' not satisfied by σ. Therefore, we have

$$\mathsf{AUT}(H; \sigma) \geq |E''| - \delta M/4 \geq (1 - 100\epsilon - 19\delta - 5\delta)m = (1 - 100\epsilon - 24\delta)m. \qquad \square$$

Proof of Claim 5.8. Let E' be the set of hyperedges in E whose vertices are all fixed points of σ. Since σ has at least $(1 - \beta)n$ fixed points, and H is $(\epsilon, 100c)$-degree bounded, by Claim 5.3, we know that $|E'| \geq m - (\epsilon + \beta) \cdot 100cn = (1 - 100(\epsilon + \beta))m$. Now consider any $e_i \in E'$, if all the edges incident to V_{C_i} are satisfied, let E'' contain i. Since there are at most $\delta M \leq 19\delta m$ edges not satisfied, we know that $|E''| \geq |E'| - 19\delta m \geq (1 - 100(\epsilon + \beta) - 19\delta)m$. We claim that for all $e_i \in E''$, the equation C_i is satisfied. Therefore we have $\mathsf{val}(\mathcal{I}; \tau) \geq 1 - 100(\epsilon + \beta) - 19\delta \geq .9 - 100(\epsilon + \beta)$, since $\delta < 1/200$.

Now we show that C_i is satisfied by τ when $e_i \in E''$. Using similar argument in the proof of Claim 5.7, one can show that when $e_i = \{j_1, j_2, j_3\} \in E''$, we have $\pi(V_{C_i}) = V_{\underline{C_i}}$. Also we have $\pi(V_{x_{j_t}}) = V_{x_{j_t}}$ for all $t \in \{1, 2, 3\}$, by the definition of E''. Let H be the induced subgraph $G_{\mathcal{C}}[V_{C_i} \cup (\cup_{t \in \{1,2,3\}} V x_{j_t})]$, let $\underline{J}$ be the induced subgraph $G_{\underline{\mathcal{C}}}[V_{\underline{C_i}} \cup (\cup_{t \in \{1,2,3\}} V_{x_{j_t}})]$. We use the following claim to conclude the proof.

Claim 5.9. *If π (after projected on the suitable vertices) is an isomorphism between J and $\underline{J}$, C_i is satisfied by τ.*

□

It remains to prove Claim 5.9. We first claim the following property about our construction $\mathcal{L}(\cdot)$.

Claim 5.10. *Let $\mathcal{C}$ be a* 3XOR *instance, let $C : x_{j_1} + x_{j_2} + x_{j_3} = b$ be an equation from $\mathcal{C}$. For any $b_1, b_2, b_3 \in \mathbb{Z}_2$, and any vertex α_C, the parity of the number of neighbors of α_C in $\{x_{j_1} \mapsto b_1, x_{j_2} \mapsto b_2, x_{j_3} \mapsto b_3\}$ is $b + b_1 + b_2 + b_3$.*

Now we are ready to prove Claim 5.9.

Proof of Claim 5.9. Suppose the equation C_i is $x_{j_1} + x_{j_2} + x_{j_3} = b$. Consider $\alpha_{C_i} \in V_{C_i}$, let $\alpha'_{\underline{C_i}} = \pi(\alpha_{C_i})$. By the construction of $G_{\mathcal{C}}$ and $G_{\underline{\mathcal{C}}}$, we know that $\alpha(x_{j_1}) + \alpha(x_{j_2}) + \alpha(x_{j_3}) = b$, and $\alpha'(x_{j_1}) + \alpha'(x_{j_2}) + \alpha'(x_{j_3}) = 0$. Now let the set $A = \{x_{j_1} \mapsto 0, x_{j_2} \mapsto 0, x_{j_3} \mapsto 0\}$. By Claim 5.10, we know that the parity of the number of neighbors of α_{C_i} in A is $b + 0 + 0 + 0 = b$. Therefore, by isomorphism, the parity of the number of neighbors of $\alpha'_{\underline{C_i}}$ in $\pi(A)$ is also b. On the other hand, by the definition of τ, we know that $\pi(A) = \{x_{j_1} \mapsto \tau(x_{j_1}), x_{j_2} \mapsto \tau(x_{j_2}), x_{j_3} \mapsto \tau(x_{j_3})\}$. By Claim 5.10 again, we know that the parity of the number of neighbors of $\alpha'_{\underline{C_i}}$ in $\pi(A)$ is $0 + \tau(x_{j_1}) + \tau(x_{j_2}) + \tau(x_{j_3})$. Therefore, we have that $\tau(x_{j_1}) + \tau(x_{j_2}) + \tau(x_{j_3}) = b$, i.e. C_i is satisfied by τ. □

6 Random graphs are robustly asymmetric

In this section we prove Theorem 1.4.

We first set up some definitions. For any graph $G = (V, E)$, let π be a permutation over the vertices in V, we write $\mathrm{id}(\pi)$ as the number of fixed points in the permutations, that is, $\mathrm{id}(\pi) = |\{v \in V : \pi(v) = v\}|$. We define $\triangle(G, \pi(G)) = \{e : e \in E, \pi(e) \notin E\} \bigcup \{e : e \notin E, \pi(e) \in E\}$. Note that $\mathsf{AUT}(G; \pi) = |E| - \frac{1}{2}|\triangle(G, \pi(G))|$.

For any permutation π over the vertex set V, we define a directed graph $G_\pi = (\binom{V}{2}, E_\pi)$, and $(e_1, e_2) \in E_\pi$ if and only if $e_2 = \pi(e_1)$. Since each $e = \{u, v\} \in \binom{V}{2}$ has in-degree and out-degree exactly 1, we can divide G_π into disjoint unions of directed cycles. We call each directed cycle a bin, and the size of the bin is the number of elements in the cycle.

Fact 6.1. *For any size-1 bins, there are only two situations:*

- $e = \{u, v\}$ *where* u *and* v *are both fixed points of* π. *We call these bins* type-1 *size-1 bins. The number of type-1 size-1 bins is at most* $\binom{\mathrm{id}(\pi)}{2}$.
- $e = \{u, v\}$ *where* $\pi(u) = v$ *and* $\pi(v) = u$. *We call these classes* type-2 *size-1 bins. The number of type-2 size-1 bins is at most* $\frac{n - \mathrm{id}(\pi)}{2}$ *for any permutation* π.

Now let us consider $G \sim \mathcal{G}_{n,m}$ where $m = cn$. Let $Z = \binom{n}{2}$ be the number of possible edges from which we choose m edges. We first prove the following lemma.

Lemma 6.2. *Let* $\mathcal{A}$ *be the event that for any permutation* π *such that* $\mathrm{id}(\pi) = (1 - \beta)n$, *the number of the edges in* G *that fall into the bins of size* ≥ 2 *is at least* $\beta m/60$ *and at most* $2\beta Z/10^9$. *Whenever* $\beta \geq \exp(-c/6)$, *we have*

$$\Pr_{G \sim \mathcal{G}_{n,m}}[\mathcal{A}] = 1 - n^{-\omega(1)}.$$

By a union bound over all the $\beta : \beta \geq \beta^*$ (where there are at most n of them), we get the following corollary.

Corollary 6.3. *Let* $\mathcal{B}$ *be the event that for every* $\beta \geq \beta^*$, *any permutation* π *such that* $\mathrm{id}(\pi) = (1 - \beta)n$, *the number of the edges in* G *that fall into the bins of size* ≥ 2 *is at least* $\beta m/60$ *at at most* $2\beta Z/10^9$. *We have*

$$\Pr_{G \sim \mathcal{G}_{n,m}}[\mathcal{B}] = 1 - n^{-\omega(1)}.$$

Before we prove Lemma 6.2, we need the following lemma.

Lemma 6.4. *Let* $G \sim \mathcal{G}_{n,m}$. *Suppose that* $\beta \geq \beta^*$. *With probability* $1 - n^{-\omega(1)}$, *for any* $T \subseteq V$, $|T| = \beta n$, *the number of edges incident to* T *is at least least* $c\beta n/40$.

Now we prove this lemma.

For any vertex $v \in V$, its expected degree in $\mathcal{G}_{n,m}$ is $2c$. We would like to prove that the probability that the degree is at most $c/10$ is very low. Indeed, we claim a more general statement.

Claim 6.5. *Let* W *be a set of* w *possible edges from* $\binom{V}{2}$, *where* $\lfloor (n-1)/2 \rfloor \leq w \leq n$,

$$\Pr_{G=(V,E) \sim \mathcal{G}_{n,m}}[|E \cap W| \leq c/10] \leq \exp(-c/2).$$

Observe that when W is the set of possible edges incident to v, Claim 6.5 says that $\Pr[\deg(v) \leq c/10] \leq \exp(-c/3)$.

For each possible edge e, we define a random variable X_e as the indicator variable for the event that e is selected as an edge in G. Claim 6.5 would be a direct application of Chernoff bound if the X_e variables were independent. However, the following claim states that Chernoff bound still holds since the variables are negatively associated. (Please refer to e.g. [JP83] for the definition of negatively associated random variables.)

Claim 6.6. *The* Z *variables* X_e *are negatively associated.*

Proof. Since $\{X_e\}$ follows the permutation distribution over m 1's and $(Z - m)$ 0's, by Theorem 2.11 in [JP83], the claim holds. □

Proof of Claim 6.5. Since $\mathbf{E}[X_e] = 2c/(n-1)$ for every e, we have $\mathbf{E}[|E \cap W|] = \sum_{e \in W} \mathbf{E}[X_e] = 2c|W|/(n-1) \geq c$. The claim follows by Chernoff bound for negatively associated variables. □

Similarly we show that

Claim 6.7. *Let W be a set of possible edges, when $c \leq n/10^{10}$,*

$$\Pr_{G=(V,E)\sim\mathcal{G}_{n,m}}[|E \cap W| \geq |W|/10^9] \leq \exp(-3|W|/10^{10}).$$

Proof. Since $\mathbf{E}[X_e] = 2c/(n-1)$ for every e, we have $\mathbf{E}[|E \cap W|] = \sum_{e \in W} \mathbf{E}[X_e] = 2c|W|/(n-1) < 3|W|/10^{10}$. The claim follows by Chernoff bound for negatively associated variables. □

Now we are ready to prove Lemma 6.4.

Proof of Lemma 6.4. Suppose the vertices of the random graph $G \sim \mathcal{G}_{n,m}$ are numbered from 1 to n. Let $X_i = \sum_{j=i+1}^{i+\lfloor (n-1)/2 \rfloor} X_{e=\{i,(j-1) \mod n+1\}}$. By Claim 6.5, we know that

$$\Pr[X_i \leq c/10] \leq \exp(-c/2).$$

Since the random variables $\{X_i\}$ are sums of disjoint sets of negatively associated random variables X_e's, we know that the X_i's are also negatively associated. Let U be the set of vertices i such that $X_i \leq c/10$. We have $\mathbf{E}[|U|] \leq n \cdot \exp(-c/2)$. Using Chernoff bound for negatively associated random variables, we have

$$\Pr\left[|U| \geq \frac{1}{2}\beta n\right] \leq \exp\left(-\frac{1}{3} \cdot \left(\frac{\beta n/2}{\mathbf{E}[|U|]} - 1\right)^2 \mathbf{E}[|U|]\right) = \exp\left(-\frac{1}{3} \cdot \left(\frac{\beta n/2}{\mathbf{E}[|U|]} - 1\right)\left(\frac{\beta n}{2} - \mathbf{E}[|U|]\right)\right).$$

Using $\beta \geq \exp(-c/6)$ and $c \geq 10$, we have $\beta n/2 - \mathbf{E}[|U|] \geq \beta n/4$. Therefore,

$$\begin{aligned}\Pr\left[|U| \geq \frac{1}{2}\beta n\right] &\leq \exp\left(-\frac{1}{3}\left(\frac{\beta}{2} \cdot \exp(c/2) - 1\right) \cdot \frac{\beta n}{4}\right) \\ &\leq \exp\left(-\frac{1}{3}\left(\frac{1}{2} \cdot \exp(c/3) - 1\right) \cdot \frac{\beta n}{4}\right) \leq \exp\left(-\frac{1}{3} \cdot \frac{1}{4} \cdot \exp(c/6) \cdot \frac{\beta n}{4}\right) \leq \exp\left(-\frac{n}{48}\right) = n^{-\omega(1)},\end{aligned}$$

where the second and fourth inequalities are because of $\beta \geq \exp(-c/6)$, and the third inequality is because of $\exp(c/3)/2 - 1 \geq \exp(c/6)/4$ for $c \geq 10$.

Therefore, with probability $1 - n^{-\omega(1)}$, there are at most $\beta n/2$ vertices with degree at most $c/10$ (since $\deg(i) \geq X_i$ for every vertex i). When this happens, for any $T \subseteq V$, $|T| = \beta \cdot n$, there are at least $(|T| - \beta n/2)$ vertices in T with degree at least $c/10$, the sum of degrees of vertices in T is at least $(|T| - \beta n/2) \cdot c/10 = c\beta n/20$, which means the number of edges incident to any vertex in T is at least $c\beta n/40$. □

Proof of Lemma 6.2. By Lemma 6.4, we know that with probability $(1 - n^{-\omega(1)})$, the number of edges in G that is incident to T is at least $c\beta n/40$, for every $T \subseteq V$ and $|T| = \beta n$. Therefore, for any π with $\mathrm{id}(\pi) = (1-\beta)n$, let T^* be the non-fixed points of π. We have $|T^*| = \beta n$. As the number of edges in size-1 bins which are incident to T is at most $|T^*|/2 = \beta n/2$, the number of selected edges that in bins of size ≥ 2 is at least $c\beta n/40 - \beta n/2 \geq \beta m/60$, when $c \geq 100$.

We also need to show that with probability $(1 - n^{-\omega(1)})$, the number of selected edges in bins of size ≥ 2 (denote this number by random variable X) is at most $2\beta^* Z/10^9$. For every π such that $\mathrm{id}(\pi) = (1-\beta)n$, let W be the set of possible edges whose end vertices are not both fixed point of π. We have $|W| = Z - \binom{(1-\beta)n}{2}$, therefore $\beta Z \leq W \leq 2\beta Z$ (for large enough n) and $X \leq |E \cap W|$. By Claim 6.7 we have $\Pr[|E \cap W| \geq |W|/10^9] \leq \exp(-3|W|/10^{10}) \leq \exp(-3 \cdot \beta Z/10^{10})$. Therefore $\Pr[X \geq 2\beta Z/10^9] \leq \exp(-3 \cdot \beta Z/10^{10}) \leq n^{-\omega(\beta n)}$. We conclude the proof by taking a union bound over $\binom{n}{\beta n} < n^{\beta n}$ ways of choosing non-fixed points for π. □

Lemma 6.8. *Conditioned on event $\mathcal{B}$, for $10^4 \leq c \leq n/10^7$, $\beta_0 \geq \beta^*$, with probability $(1 - n^{-17})$, G is $(\beta_0, \beta_0/240)$-asymmetric.*

Proof. For any permutation π, we define a set of more fine-grained bins. We start with the bins we defined before, and split the bins of size ≥ 4 into bins of size 2 and at most one bin of size 3 as follows. Suppose the original bin contains $\{e_1, e_2, \ldots, e_l\}$, where $\pi(e_i) = e_{i+1}$ and $\pi(e_l) = e_1$, $l \geq 4$, the we have new bins which contains $\{e_1, e_2\}, \ldots, \{e_{l-1}, e_l\}$ if l is even, and $\{e_1, e_2\}, \ldots, \{e_{l-4}, e_{l-3}\}, \{e_{l-2}, e_{l-1}, e_l\}$ if l is odd.

For each bin of size 2 and size 3, if all the edges are in G, we call it a full bin; if none of them are in G, we call it an empty bin; otherwise, we call it a half-full bin. Fix a permutation π, let s_π be the number of half-full bins. For each half-full bin, it contributes at least one to $\triangle(G, \pi(G))$, therefore $s_\pi \leq |\triangle(G, \pi(G))| = 2(m - \mathsf{AUT}(G, \pi))$. We have

$$\begin{aligned}\Pr[G \text{ is } (\beta_0, \beta_0/240)\text{-asymmetric}|\mathcal{B}] &= \Pr[\forall \pi : \mathrm{id}(\pi) \leq (1-\beta_0)n, \mathsf{AUT}(G,\pi) \leq (1-\beta_0/240)m|\mathcal{B}] \\ &\geq \Pr[\forall \pi : \mathrm{id}(\pi) \leq (1-\beta_0)n, s_\pi \geq \beta_0 m/120|\mathcal{B}].\end{aligned}$$

Now we turn to show that

$$\Pr[\forall \pi : \mathrm{id}(\pi) \leq (1-\beta_0)n, s_\pi \geq \beta_0 m/120|\mathcal{B}] \geq 1 - n^{-17}.$$

To show this, we only have to prove for every β where $\beta \geq \beta_0$,

$$\Pr[\forall \pi : \mathrm{id}(\pi) = (1-\beta)n, s_\pi \geq \beta m/120|\mathcal{B}] \geq 1 - n^{-18}, \tag{4}$$

and take a union bound over (at most n possible) β's.

Fix $\beta : \beta \geq \beta_0$ and fix a permutation π such that $\mathrm{id}(\pi) = (1-\beta)n$. Let $\mathcal{C}_t$ be the event that $\mathcal{B}$ happens and there are t edges in G fall into the bins of size 2 and 3. Since $\mathcal{B}$ is a disjoint union of $\mathcal{C}_t$ for all $t : \beta m/60 \leq t \leq 2\beta Z/10^9$, we have

$$\Pr[s_\pi \leq \beta m/120|\mathcal{B}] = \sum_{t=\beta m/60}^{m} \Pr[s_\pi \leq \beta m/120|\mathcal{C}_t] \cdot \Pr[\mathcal{C}_t|\mathcal{B}].$$

We will prove that

$$\Pr[s_\pi \leq \beta m/120|\mathcal{C}_t] = \left(\binom{n}{\beta n}(\beta n)!\right)^{-1} \cdot n^{-18}, \tag{5}$$

and by taking a union bound over all $\binom{n}{\beta n}(\beta n)!$ possible π's, we prove (4).

Let $\gamma = \beta/120$. Let L be the number of possible edges in bins of size 2 and 3. We have $L = Z - \binom{(1-\beta)n}{2} - \frac{\beta n}{2} \geq \beta n^2/4$ (for large enough n). Together with $t \leq 2\beta Z/10^9$, we have $t \leq 4L/10^9 \leq L/10^8$. Let B be the number of bins of size 2 and 3. Fix t such that $\beta m/60 \leq t \leq L/10^8$. Conditioned on $\mathcal{C}_t$, the $\binom{L}{t}$ ways to select these t edges are uniformly distributed. Now we compute the number of ways such that there are at most $2\gamma m$ half-full bins. Suppose that there are i half-full bins (for $i \leq 2\gamma m \leq t/2$). There are $\binom{B}{i}$ ways to choose these bins. There are at most $\binom{B}{(t-i)/2}$ ways to choose the full bins (since $t/2 < L/2 \cdot 10^8 < B/2$). For each half-full bin, there are at most 6 ways to choose the edges in the bin. Therefore,

$$\Pr[s_\pi = i|\mathcal{C}_t] \leq 6^i \binom{B}{i}\binom{B}{\frac{t-i}{2}}\binom{L}{t}^{-1} \leq 6^i \cdot \frac{L^{\frac{t+i}{2}}}{i!\left(\frac{t-i}{2}\right)!}\left(\frac{t}{L}\right)^t = 6^i L^{-\frac{t-i}{2}} \frac{t^t}{i!\left(\frac{t-i}{2}\right)!}. \tag{6}$$

Since $i! \geq (i/e)^i$, we have

$$(6) \leq 6^i L^{-\frac{t-i}{2}} e^{\frac{t+i}{2}} \frac{t^t}{i^i\left(\frac{t-i}{2}\right)^{\frac{t-i}{2}}} \leq 6^i L^{-\frac{t-i}{2}} (2e)^{\frac{t+i}{2}} \frac{t^t}{i^i (t-i)^{\frac{t-i}{2}}}. \tag{7}$$

Using $i^i(t-i)^{\frac{t-i}{2}} \geq \left(\frac{t}{4}\right)^{\frac{t+i}{2}}$, we have

$$(7) \leq 6^i (8e)^{\frac{t+i}{2}} \left(\frac{t}{L}\right)^{\frac{t-i}{2}} \leq (48e)^{\frac{3t}{4}} \left(\frac{t}{L}\right)^{\frac{t}{4}} \leq \left(\frac{10^7 t}{L}\right)^{\frac{t}{4}}. \tag{8}$$

Since $\beta m/60 \le t \le L/10^8$, and $(10^7 t/L)^{t/4}$ is monotonically decreasing when $t < L/(10^7 e)$, we have

$$(8) \le \left(\frac{10^7\beta m}{60L}\right)^{\frac{\beta m}{240}} \le \left(\frac{10^7\beta cn}{60\beta n^2/4}\right)^{\frac{\beta cn}{240}} \le \left(\frac{10^6 c}{n}\right)^{\frac{\beta cn}{240}} \le \left(\frac{10^{10}}{n}\right)^{40\beta n} \le \left(\frac{1}{n}\right)^{30\beta n},$$

where the second last inequality is because $10^4 \le c \le n/10^7$ and $(10^6 c/n)^{\beta cn/240}$ is monotonically decreasing in this range, and the last inequality is for large enough n. Observing that $\binom{n}{\beta n}(\beta n)! \le n^{2\beta n}$ and $\beta \ge \beta^* \ge 1/n$, we proved that

$$\Pr[s_\pi = i | \mathcal{C}_t] \le \left(\binom{n}{\beta n}(\beta n)!\right)^{-1} \cdot n^{-20}.$$

By taking a union bound over all (at most n^2 many) i's such that $i \le 2\gamma m$, we prove (5). □

Theorem 1.4 is proved by combining Corollary 6.3, Lemma 6.8, and taking a union bound over all possible $\beta = \beta_0$ (where there are at most n of them).

6.1 Generalization to hypergraphs

In this subsection, we generalize Theorem 1.4 to random k-uniform hypergraphs for any constant $k \ge 3$.

Theorem 6.9. *For any constant k, there exists constant κ_k, such that for $m = cn$ where $\kappa_k \le c \le \binom{n}{k}/\kappa_k^3$, n large enough, if we set $\beta^* = \max\{\exp(-c/6), 1/n\}$, with probability $(1 - n^{-15})$, for all $\beta : \beta^* \le \beta \le 1$, a random graph H from the distribution $\mathcal{G}^{(k)}_{n,m}$ is $(\beta, \beta/240)$-asymmetric. For $k = 3$, $\kappa_3 = 10^4$ suffices.*

The proof of Theorem 6.9 mostly follows the lines of the proof of Theorem 1.4. But we need some small modifications. For simplicity, we only prove the theorem for $k = 3$. For higher k, we encourage the readers to check by themselves.

Now we work with $k = 3$. For any permutation π over the vertex set V, we define a directed graph $G^{(3)}_\pi = (\binom{V}{3}, E_\pi)$, and $(e_1, e_2) \in E_\pi$ if and only if $e_2 = \pi(e_1)$. Since each $e \in \binom{V}{3}$ has in-degree and out-degree exactly 1, we can divide $G^{(3)}_\pi$ into disjoint unions of directed cycles. Similarly as in the ordinary graph case, we call each directed cycle a bin, and the size of the bin is the number of elements in the cycle. Let $Z_3 = \binom{n}{3}$.

Fact 6.10. *For any size-1 bins, there are only three situations:*

- *$e = \{u, v, w\}$ where u, v and w are all fixed points of π. We call these bins* type-1 *size-1 bins.The number of type-1 size-1 bins is at most $\binom{\mathrm{id}(\pi)}{3}$.*
- *$e = \{u, v, w\}$ where one of them is a fixed point of π and the other two map to each other under π. We call these classes type-2 size-1 bins. The number of type-2 size-1 bins is at most $\mathrm{id}(\pi) \cdot \frac{n-\mathrm{id}(\pi)}{2} = O(n^2)$ for any permutation π.*
- *$e = \{u, v, w\}$ where $\pi(u) = v$ and $\pi(v) = w$ and $\pi(w) = u$. We call these classes type-3 size-1 bins. The number of type-3 size-1 bins is at most $\frac{n-\mathrm{id}(\pi)}{3}$ for any permutation π.*

The following lemma is an analogue of Lemma 6.2.

Lemma 6.11. *For any fixed β where $\beta \ge \beta^*$, let $\mathcal{D}'$ be the event that for any permutation π such that $\mathrm{id}(\pi) = (1 - \beta)n$, the number of the hyperedges in H that fall into the bins of size ≥ 2 is at least $\beta m/60$ and at most $2\beta Z_3/10^9$. We have*

$$\Pr_{H \sim \mathcal{H}^{(3)}_{n,m}}[\mathcal{D}'] = 1 - n^{-\omega(1)}.$$

By a union bound over all the $\beta : \beta \ge \beta^*$ (there are at most n of them), we get the following corollary.

Corollary 6.12. *Let $\mathcal{D}$ be the event that for every $\beta \ge \beta^*$, any permutation π such that $\mathrm{id}(\pi) = (1 - \beta)n$, the number of the hyperedges in H that fall into the bins of size ≥ 2 is at least $\beta m/60$ at at most $2\beta Z_3/10^9$. We have:*

$$\Pr_{H \sim \mathcal{H}^{(3)}_{n,m}}[\mathcal{D}] = 1 - n^{-\omega(1)}.$$

The proof of Lemma 6.11 is similar to that of Lemma 6.2, except that now we also need to take care of type-2 size-1 bins.

Lemma 6.13. *With probability $1 - n^{\omega(1)}$, for any permutation π with $\mathrm{id}(\pi) \le (1-\beta^*)n$, the number of selected hyperedges in type-2 size-1 bins is at most $c\beta n/1000$.*

Proof. We first prove that for every $\beta \ge \beta^*$, with probability $1 - n^{-\omega(1)}$, for every permutation π with $\mathrm{id}(\pi) = (1-\beta)n$, the number of selected hyperedges in type-2 size-1 bins is at most $c\beta n/1000$. By a union bound over all possible β's (where there are at most n of them), we get the desired statement.

For any fixed permutation π with $\mathrm{id}(\pi) = (1-\beta)n$, the number of type-2 size-1 bins is at most $(1-\beta)n\beta n/2 \le \beta n^2/2$. For each possible hyperedge e, we define the random variable X_e as the indicator variable for the event that e is selected as an hyperedge in H. Note that $\mathbf{E}[X_e] = cn/Z_3 \le \frac{7c}{n^2}$. Define random variable $X = \sum_{e \text{ in type-2 size-1 bin}} X_e$ as the number of selected hyperedges in type-2 size-1 bins, by linearity of expectation,

$$\mathbf{E}[X] \le \frac{\beta n^2}{2} \cdot \frac{7c}{n^2} < 4c\beta.$$

On the other hand, we can also show that all these random variables are negative associated, therefore through Chernoff bound for negative associated random variables, we have

$$\Pr_{H \sim \mathcal{G}^{(3)}_{n,m}}[X \ge \beta cn/1000] \le \exp(-1/3 \cdot (n/250 - 1)^2 \cdot 4c \cdot \beta) \le \exp(-c\beta n^2/10^5)$$

By a union bound over at most $\binom{n}{\beta n}(\beta n)! \le n^{2\beta n}$ such permutations, the probability that there exists π with $\mathrm{id}(\pi) = (1-\beta)n$ such that the number of type-2 size-1 bins is more than $c\beta n$, is at most

$$\exp(-c\beta n^2/10^5) \cdot n^{2\beta n} = \exp(-c\beta n^2/(10^5) + 2\beta n \log n) \le \exp(-c\beta n^2/(10^6)) \le n^{-\omega(1)}.$$

□

The following lemma is an analogue of Lemma 6.4, and the proof is almost identical.

Lemma 6.14. *Let $H \sim \mathcal{G}^{(3)}_{n,m}$. Suppose that $\beta \ge \beta^*$. With probability $1 - n^{-\omega(1)}$, for any $T \subseteq V$, $|T| = \beta n$, the number of hyperedges incident to T is at least least $c\beta n/40$.*

Proof of Lemma 6.11. We only establish the lower bound ($\beta m/60$). The proof for upper bound is almost identical to that in Lemma 6.2.

Let T be the set of non-fixed points of π, then $|T| = \beta n$. By Lemma 6.14, we know that with probability $(1 - n^{-\omega(1)})$, there are at least $c\beta n/40$ hyperedges incident to T – all these edges are either in bins of size ≥ 2, or edges in size-1 bins of type-2 or type-3. By Lemma 6.13, we know that with probability $1 - n^{\omega(1)}$ the number of selected hyperedges that fall into type-2 size-1 bins is at most $c\beta n/1000$. Finally we recall that there are at most $\frac{n - \mathrm{id}(\pi)}{3} = \beta n/3$ type-3 size-1 bins.

Therefore, with probability $(1 - n^{-\omega(1)}$, the number of selected hyperedges that fall into bins of size ≥ 2 is at least $c\beta n/40 - \beta n/3 - c\beta n/1000 \ge c\beta n/60$. □

Finally we state the following analogue of Lemma 6.8 (whose proof is also almost identical).

Lemma 6.15. *Conditioned on event $\mathcal{D}$, for $10^4 \le c \le n/10^7$, $\beta \ge \beta^*$, with probability $(1 - n^{-17})$, $H \sim \mathcal{G}^{(3)}_{n,m}$ is $(\beta, \beta/240)$-asymmetric.*

The $k = 3$ case in Theorem 6.9 follows from Corollary 6.12 and Lemma 6.15.

7 Generalizing Feige's Hypothesis

In this section, we show how our main results can be derived from a much broader class of assumptions than just Feige's R3XOR hypothesis. In doing so, we are inspired by the recent NP-hardness results of [Cha13], along with the Lasserre gaps of [Tul09]. The focus of these papers was on a specific type of predicates, which we now define.

Definition 7. *Let $k \geq 3$, and let H be a finite abelian group. Then a predicate $\psi : H^k \to \{0,1\}$ is a* balanced pairwise independent subgroup *of H^k if the following two conditions hold.*

- *For a uniformly random element $a \sim \psi$, each coordinate a_i is uniformly distributed over G. Furthermore, each pair of coordinates a_i and a_j, for $i \neq j$, is independent.*
- *ψ is a subgroup of H^k.*

Here we are using the convention that ψ is both a function on H^k and a subset of H^k, i.e. the subset $\psi^{-1}(1)$. We will everywhere assume that ψ is a proper subgroup of H^k, as otherwise ψ is uninteresting.

An instance $\mathcal{C}$ of the Additive$-$CSP(ψ) problem consists of a set of constraints $\mathcal{C} = (C_1, \ldots, C_m)$, each of the form

$$\psi(x_{j_1} + a_1, \ldots, x_{j_k} + a_k) = 1,$$

where the a_i's are elements of H. The corresponding homogeneous instance $\underline{\mathcal{C}} = \{\underline{C_1}, \ldots, \underline{C_m}\}$ is formed by setting all of the a_i's in each constraint to 0. We note that because ψ is a subgroup of H^k, it contains 0^k, so the assignment $x_1 = \ldots = x_n = 0$ satisfies every constraint in $\underline{\mathcal{C}}$.

For a fixed number of variables n and constraints m, a random instance of Additive$-$CSP(ψ) is formed by generating m constraints of this form independently and uniformly. This involves, for each constraint, picking the variables $x_{j_1}, \ldots, x_{j_k}$ uniformly from the $\binom{n}{k} \cdot k!$ possibilities, along with choosing each a_i independently and uniformly at random. The next hypothesis states that random Additive$-$CSP(ψ) instances are indistinguishable from almost-satisfiable instances.

Random Additive$-$CSP(ψ) **Hypothesis.** *For every fixed $\epsilon > 0$, $\Delta \in Z^+$, there is no polynomial time algorithm which on almost all* Additive$-$CSP(ψ) *instances with n variables and $m = \Delta n$ constraints outputs "typical", but which never outputs "typical" on instances which an assignment satisfying at least $(1-\epsilon)m$ constraints.*

Note that Feige's R3XOR hypothesis can be recovered by setting $\psi = 3\text{XOR}$. Using this, we can weaken the assumptions for Theorem 1.2 so that it only needs the Random Additive$-$CSP(ψ) Hypothesis to be true for some balanced pairwise independent subgroup ψ.

Theorem 7.1. *Let ψ be a predicate which is a balanced pairwise independent subgroup. Assume the Random* Additive$-$CSP(ψ) *Hypothesis. Then there is no polynomial-time algorithm for* RobustGiso. *More precisely, there is an absolute constant $\delta > 0$ such that the following holds: suppose there is a $t(n)$-time algorithm which can distinguish $(1-\epsilon)$-isomorphic n-vertex, m-edge graph pairs from pairs which are not even $(1-\delta)$-isomorphic (where $m = O(n)$). Then there is a universal constant $\Delta \in \mathbb{Z}^+$ and a $t(O(n))$-time algorithm which outputs "typical" for almost all n-variable, Δn-constraint instances of the* Additive$-$CSP(ψ) *problem, yet which never outputs "typical" on instances which are $(1-\epsilon)$-satisfiable.*

Outline of proof. The proof of Theorem Theorem 7.1 largely follows the same outline as the proof of Theorem 1.2. As a result, our proof of Theorem 7.1 will be carried out at a somewhat high level, except in those areas where we need to highlight changes. The most significant change is that the general abelian group setting requires us to develop more complicated gadgets, and we will need to describe these gadgets before stating the reduction.

7.1 Abelian variable gadgets

For a given variable x, the construction in Section 3 creates two vertices $x \mapsto 0$ and $x \mapsto 1$ and places an edge between them. This subgraph has two automorphisms, corresponding to the two possible Boolean assignments

to x. This is a simple example of what we will call a *variable gadget*. When x takes values from a more general group H, a variable gadget for x will be a graph $G = (V_x \cup V_A, E)$. Here, V_x is the set of vertices "$x \mapsto a$" for $a \in H$, and the vertices in V_A are thought of as auxiliary vertices. Furthermore, we require that the only automorphisms of G correspond to the $|H|$ possible assignments to x. More formally:

Definition 8. *Let x be a variable which takes values from H, an abelian group. A* variable gadget *for x is a graph $G = (V_x \cup V_A, E)$ satisfying*

> ***(Completeness)*** *For each $a \in H$, G has an automorphism f_a such that $f_a(x \mapsto b) = x \mapsto (a + b)$ for each $b \in H$.*
>
> ***(Soundness)*** *Let $f : V_x \cup V_A \to V_x \cup V_A$ be an automorphism on G. Then there exists an element $a \in H$ such that $f(x \mapsto b) = x \mapsto (a + b)$ for each $b \in H$.*

Constructing a variable gadget for x when x is non-Boolean is more complicated than just slapping an edge between two vertices. However, we are able to construct variable gadgets with the following parameters:

Lemma 7.2. *Let x be a variable over H, an abelian group. Then there is a variable gadget for x with at most $4 \cdot |H| \log^2 |H|$ auxiliary variables, $7 \cdot |H| \log^2 |H|$ edges, and a maximum clique size of $\min\{4, |H|\}$.*

The proof of this lemma can be found in Section 7.5.

7.2 Label-extended graphs

Let $\psi : H^k \to \{0, 1\}$ be a subgroup, and let C be an $\mathsf{Additive{-}CSP}(\psi)$ constraint. The *label-extended graph* of C consists of vertex sets V_{x_i} as defined in Section 7.1 for each $i \in [k]$; and $|\psi|$ "constraint vertices" with names "$x_1 \mapsto a_1, \ldots, x_k \mapsto a_k$" for each assignment $(a_1, \ldots, a_k)$ which satisfies C. We will sometimes write these constraint vertices as assignments $\alpha : \{x_1, \ldots, x_k\} \to H$, and the set of them we call V_C. For edges, each constraint vertex α is connected to each of the k variable vertices it is consistent with.

Let G be the label-extended graph of a constraint C, and let $\underline{G}$ be the label-extended graph of the homogeneous constraint $\underline{C}$. Our goal is to show that every homomorphism from G to $\underline{G}$ corresponds to a satisfying assignment of C. While this may not be true in general, we will be able to restrict our attention to only those automorphisms which act on each set V_{x_i} in a very limited way; this is due to the variable gadget from Section 7.1. Let $\alpha : \{x_1, \ldots, x_k\} \to H$ be an assignment to the variables. Then a permutation f on the vertices of G is an α-*assignment* if for each $i \in [k]$, $f(x_i \mapsto b) = x_i \mapsto (b - \alpha(x_i))$ for each $b \in H$. Then

Lemma 7.3. *Let $\psi : H^k \to \{0, 1\}$, let C be an $\mathsf{Additive{-}CSP}(\psi)$ constraint, and let G be its label-extended graph. Let $\underline{G}$ be the label extended graph of $\underline{C}$. Then*

> ***(Completeness)*** *For each assignment α which satisfies C, there is an α-permutation f which is a homomorphism from G to $\underline{G}$.*
>
> ***(Soundness)*** *Let f be an α-permutation which is a homomorphism from G to $\underline{G}$. Then α is a satisfying assignment of C.*

In fact, though we will not use this, it can be shown that the soundness holds even if f is only a $(1 - \frac{1}{k} + \epsilon)$ homomorphism from G to $\underline{G}$, for any $\epsilon > 0$. Lemma 7.3 is proved in Section 7.6.

7.3 Reduction from $\mathsf{Additive{-}CSP}(\psi)$ to GISO

We begin with a generalization of Definition 4:

Definition 9. *Let C be a ψ constraint involving variables $x_1, \ldots, x_k$. The associated gadget graph G_C has a vertex set consisting of the variable vertices V_{x_i} for each $i \in [k]$ and the constraint vertices V_C. For each $i \in [k]$, the variable gadget from Lemma 7.2 is added to V_{x_i}; this adds a set of auxiliary variables, which we call V_{A_i}. Next, V_C is connected to the variable vertex sets via the label-extended graph. Finally, the variables in V_C are connected to each other by a clique.*

Here the vertex sets V_{x_i} are as defined in Section 7.1 and the vertex set V_C is as defined in Section 7.6.

Let $\mathcal{C}$ be a collection of Additive$-$CSP(ψ) constraints. We define the associated graph $G_{\mathcal{C}}$ as in Definition 5 as a union of gadget graphs G_C where the variable vertex sets and gadgets are associated with each other.

The number variable vertices in $G_{\mathcal{C}}$ is $|H| \cdot n$, the number of constraint vertices is $|\psi| \cdot m$, and the number of auxiliary variables given in Lemma 7.2 is at most $4 \cdot |H| \log^2 |H| \cdot n$. Next, the number of edges from the complete graphs and the constraint-variable graphs is $\left(\binom{|\psi|}{2} + |\psi| \cdot k\right)$, and the number of edges given by Lemma 7.2 is at most $7 \cdot |H| \log^2 |H| n$. This gives us the following remark:

Remark 7.1. *If $\mathcal{C}$ is an* Additive$-$CSP(ψ) *instance with n vertices and m constraints then the graph $G_{\mathcal{C}}$ has $N = N_1 m + N_2 n$ vertices and $M = M_1 m + M_2 n$ edges. Here the N_i's and M_i's are absolute constants which depend only on the group H and the predicate φ.*

One thing that we will need is that the largest cliques of $G_{\mathcal{C}}$ are those only involving constraint vertices. To show this, we will first need the following simple proposition:

Proposition 7.4. *Let H be a finite abelian group and let $\psi : H^k \to \{0,1\}$ be a balanced pairwise independent subgroup. Then $|\psi| \geq |H|^2$.*

Proof. For a uniformly random element $a \sim \psi$, the definition of a balanced pairwise independent subgroup states that a_1 and a_2 will each be uniform on H and independent of each other. Thus, (a_1, a_2) ranges over the $|H|^2$ possibilities with equal probability, meaning that ψ must have at least $|H|^2$ elements. □

Remark 7.2. *By Lemma 7.2, the largest clique in the variable gadgets is of size $\min\{4, |H|\}$, whereas the constraint vertices come in cliques of size $|\psi| \geq |H|^2$. Thus, the largest cliques in the graph consist of the constraint vertices.*

The reduction. *Given a collection of* Additive$-$CSP(ψ) *constraints $\mathcal{C}$, we introduce the corresponding* GISO *instance $(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}})$.*

7.4 Abelian proofs

Having defined the reduction from Additive$-$CSP(ψ) to GISO, we are now in a position to prove Theorem 7.1. Our two main lemmas are:

Lemma 7.5. *Let $\mathcal{C}$ be a* Additive$-$CSP(ψ) *instance such that* val$(\mathcal{C}) \geq 1 - \epsilon$. *Then* GI$(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq 1 - \epsilon$.

Lemma 7.6. *There are absolute constants $\delta, c > 0$ depending only on the predicate ψ such that the following holds: let $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$ be a random* Additive$-$CSP(ψ) *instance with n variables and $m = cn$ equations. With probability $1 - o(1)$, we have*

$$\text{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) < 1 - \delta.$$

Given a polynomial-time algorithm $\mathcal{A}$ for ROBUSTGISO, we can define a polynomial-time random Additive$-$CSP(ψ) solver $\mathcal{A}_\psi$ which, on input $\mathcal{C}$, runs $\mathcal{A}$ on $(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}})$. If val$(\mathcal{C}) \geq 1 - \epsilon$, then by Lemma 7.5, GI$(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq 1 - \epsilon$, and so $\mathcal{A}_\psi$ will never output "typical". On the other hand, if $\mathcal{C}$ is a random Additive$-$CSP(ψ) instance, then with probability $1 - o(1)$ we have GI$(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) < 1 - \delta$ by Lemma 7.6, so $\mathcal{A}_\psi$ will almost always output "typical". However, this contradicts the Random Additive$-$CSP(ψ) hypothesis, meaning that no such ROBUSTGISO algorithm can exist.

We now proceed to the proofs.

Completeness This case follows directly from the completeness properties of our gadgets.

Proof of Lemma 7.5. Let τ be an assignment to the variables in $\mathcal{C}$ such that val$(\mathcal{C}; \tau) \geq 1 - \epsilon$. Now we define a bijection π from the variables in $G_{\mathcal{C}}$ to the ones in $G_{\underline{\mathcal{C}}}$.

For each variable vertex x_j, set $a_j := \tau(x_j)$. By Lemma 7.2 there is an automorphism $f_{a_j} : (V_{x_j} \cup A_{x_j}) \to (V_{x_j} \cup A_{x_j})$ such that $f(x_j \to b) = x_j \to (b - a_j)$ for each $b \in H$. Set π to be consistent with f_{a_j} on $V_{x_j} \cup A_{x_j}$ for each $j \in [n]$.

Next, let C_i be a constraint which is satisfied by τ, and assume without loss of generality that it is applied to the variables $x_1, \ldots, x_k$. Consider the subgraph G_i which is the induced subgraph of $G_{\mathcal{C}}$ restricted to the vertices in $V_{x_1}, \ldots, V_{x_k}$ and V_{C_i}. Define $\underline{G_i}$ to be the analogous subgraph of $G_{\underline{\mathcal{C}}}$. Then π restricted G_i is an $(a_1, \ldots, a_k)$-permutation, as defined in Section 7.6, and so Lemma 7.3 says that it can be extended to a homomorphism between G_i and $\underline{G_i}$. Define π on V_{C_i} accordingly.

This leaves how to define π on the variables in V_{C_i} when τ does not satisfy C_i, and here we allow π to map V_{C_i} arbitrarily into $V_{\underline{C_i}}$. Until this step, π was a homomorphism on every subgraph it was defined on. Thus, the only edges violated by π are those involved in the label-extended graph of any constraint C_i which τ does not satisfy. As these constraints make up only an ϵ-fraction of the total constraints, $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq (1-\epsilon)$. □

Soundness In this section, we sketch the proof of the soundness lemma. To start with:

Claim 7.7. *Suppose $c \geq 1$. A random k-uniform hypergraph G drawn from $\mathcal{G}^{(k)}_{n,m}$, where $m = cn$, is $(1/kc, 100kc)$-degree bounded with probability $1 - o(1)$.*

Lemma 7.6 is directly implied by the following two lemmas.

Lemma 7.8. *Let $G = ([n], E = \{e_i\})$ be a k-uniform hypergraph with n vertices and $m = cn$ hyperedges. Suppose H is $(\epsilon, 100kc)$-degree bounded and (β, γ)-asymmetric, where $\gamma \geq 200k\epsilon$. Let $\mathcal{C} = \{C_1, C_2, \ldots, C_m\}$ be an arbitrary $\mathsf{Additive{-}CSP}(\psi)$ instance with n variables and m constraints based on H. If we set*

$$\delta := \delta(c, \epsilon, \beta, \gamma) = \min\left\{\frac{1}{10(M_1+M_2)}, \frac{\gamma}{4(M_1+M_2)}, \frac{\epsilon}{3(M_1+M_2)N_1 c}\right\},$$

when $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq 1 - \delta$, we have $\mathsf{val}(\mathcal{C}) \geq .9 - 100(\epsilon + \beta)$.

Lemma 7.9. *Given a predicate ψ, there is a constant $c > 0$ such that the following holds: Let $\mathcal{C}$ be a random $\mathsf{Additive{-}CSP}(\psi)$ instance with n and $m \geq cn$ equations. With probability $1 - o(1)$, we have $\mathsf{val}(\mathcal{C}) < 0.51$.*

Proof of Lemma 7.6. We will choose a value for $c \geq 10^{10}$ large enough so that our applications of Lemma 7.9 and Theorem 6.9 work properly. Set $\epsilon = \frac{1}{kc}$, $\gamma = \frac{200}{c}$, and $\beta = \frac{48000}{c}$. Combining Claim 7.7, Lemma 7.9, and Theorem 6.9, we know that with probability $(1 - o(1))$, all of the following hold:

1. G is $(\epsilon, 100kc)$-degree bounded,
2. G is (β, γ)-asymmetric, and
3. $\mathsf{val}(\mathcal{C}) < 0.51$.

Given that these hold, assume for sake of contradiction that $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) \geq 1 - \delta$, for δ as defined in Lemma 7.9. Then because G satisfies Item 1 and Item 2, Lemma 7.9 implies that

$$\mathsf{val}(\mathcal{C}) \geq .9 - 100\left(\frac{1}{kc} + \frac{48000}{c}\right) \geq .8,$$

where the last step follows because $c \geq 10^{10}$. However, this contradicts Item 3. Therefore, $\mathsf{GI}(G_{\mathcal{C}}, G_{\underline{\mathcal{C}}}) < 1 - \delta$ with probability $1 - o(1)$. □

The proof of Lemma 7.9 is standard.

Proof of Lemma 7.9. First, we note that because ψ is a proper subgroup of H^k, $|\psi|/|H^k| \leq 1/|H| \leq 1/2$. Thus, the probability that a random assignment satisfies ψ is at most $1/2$. Fix an assignment α to the n variables, and let $\mathcal{C}$ be a random $\mathsf{Additive{-}CSP}(\psi)$ instance. By the Chernoff bound,

$$\Pr[\alpha \text{ satisfies at least } .51m \text{ constraints}] \leq e^{-\Theta(m)} = e^{-\Theta(c \cdot n)}.$$

Union bounding over all the $|H|^n$ possible assignments, the probability that $\mathsf{val}(\mathcal{C}) \geq .51m$ is at most $|H|^n e^{-\Theta(c \cdot n)}$. Thus, by taking $c = \Theta(\log |H|)$ this probability is $o(1)$. □

Finally, we give a sketch of the main lemma.

Sketch of Lemma 7.8. Let π be a bijection mapping the vertices in $\mathcal{L}(\mathcal{C})$ to the vertices $\mathcal{L}(\underline{\mathcal{C}})$ such that $\mathsf{GI}(\mathcal{L}(\mathcal{C}), \mathcal{L}(\underline{\mathcal{C}}); \pi) \geq 1 - \delta$. Let A be the set of $i \in [m]$ such that $\pi(V_{C_i}) = V_{\underline{C_{i'}}}$ for some i', and let B be the set of $j \in [n]$ such that $\pi(V_{x_j} \cup V_{A_j}) = V_{x_{j'}} \cup V_{A_{j'}}$ for some j'. Then

Claim 7.10. $|A| \geq (1 - \delta(M_1 + M_2))m$, $|B| \geq (1 - 2\delta(M_1 + M_2)N_1 c)n$.

The proof of this statement is essentially identical to the proof of Claim 5.6. The only difference here is that we use 7.2 to show that the unique largest cliques are the constraint vertices. Now, we would like to show that $\pi(V_{x_j})$ acts in a highly structured way for most of the $j \in B$, as in the definition of a variable gadget. In particular, let B_{shift} be the set of $j \in [n]$ such that for some j' and $a \in H$, $\pi(x_j \mapsto b) = (x'_j \mapsto (b - a))$ for all $b \in H$. Then

Claim 7.11. $|B_{shift}| \geq (1 - 3\delta(M_1 + M_2)N_1 c)n$.

Proof. Consider $B'_{shift} = B \cap B_{shift}$. Let $j \in B$ and suppose $x_{j'}$ is the coordinate π maps x_j to, i.e. $\pi(V_{x_j} \cup V_{A_j}) = V_{x_{j'}} \cup V_{A_{j'}}$. If π is *not* of the form $\pi(x_j \mapsto b) = x_{j'} \mapsto (b - a)$, then the definition of a variable gadget tells us that π cannot be a homomorphism on the induced subgraphs on $V_{x_j} \cup V_{A_j}$ and $V_{x_{j'}} \cup V_{A_{j'}}$. In other words, π must not satisfy one of the edges in $V_{x_j} \cup V_{A_j}$. Thus, $|B \setminus B'_{shift}| \leq \delta M \leq \delta(\beta_1 + \beta_2)m = \delta(\beta_1 + \beta_2)cn$. Subtracting this from the size of B gives the claim. □

Now, define a permutation σ on the variables in $\mathcal{C}$ such that $\sigma(j) = j'$ whenever $\pi(V_{x_j}) = V_{x_{j'}}$ and $j \in B_{shift}$. Furthermore, ensure that $\sigma(j) \neq j$ for any $j \notin B_{shift}$. Then σ is an almost-automorphism for the hypergraph $G = ([n], E)$:

Claim 7.12. $\mathsf{AUT}(G; \sigma) \geq (1 - 100k\epsilon - 2\delta(M_1 + M_2))m$.

By our setting of δ, we have $2\delta(M_1 + M_2) \leq \gamma/2$, and we also know that $100k\epsilon \leq \gamma/2$. Thus, $\mathsf{AUT}(G, \sigma) \geq (1 - \gamma)m$, and therefore σ has at least $(1 - \beta)n$ fixed points.

We can now define an assignment $\tau : \{x_j\} \to H$. For each j which is not a fixed point of σ, define $\tau(x_j)$ arbitrarily. For each j which is a fixed point of σ, we ensured when constructing σ that $j \in B_{shift}$. Thus, there is some $a \in H$ such that $\pi(x_j \mapsto b) = (x_j \mapsto (b - a))$ for all $b \in H$. Using this, define $\tau(x_j) = a$. We will now show the following claim:

Claim 7.13. $\mathsf{val}(\mathcal{C}, \tau) \geq .9 - 100(\epsilon + \beta)$.

Proof. Let E' be the set of hyperedges in E whose vertices are all fixed points of σ, and let $E'' \subset E'$ contain those hyperedges e_i in E' for which all the edges incident to V_{C_i} are satisfied. By an argument similar to Claim 5.8, we know that $|E''| \geq (1 - 100(\epsilon + \beta) - (M_1 + M_2)\delta)m \geq .9 - 100(\epsilon + \beta)$, by our choice of δ.

Given $e_i \in E''$, assume without loss of generality that $e_i = \{x_1, \ldots, x_k\}$. Then we know that for some $a_1, \ldots, a_k \in H$, $\pi(x_j \mapsto b) = x_j \mapsto (b_j - a_j)$ for each j. Consider the subgraph G_i which is the induced subgraph of $G_{\mathcal{C}}$ restricted to the vertices in $V_{x_1}, \ldots, V_{x_k}$ and V_{C_i}. Define $\underline{G_i}$ to be the analogous subgraph of $G_{\underline{C}}$. Then π restricted to G_i is an $(a_1, \ldots, a_k)$-permutation, and it is a homomorphism from G_i to $\underline{G_i}$. Therefore, Lemma 7.3 tells us that $(a_1, \ldots, a_k)$ is a satisfying assignment of C_i. Thus, for each $T_i \in T''$, the assignment τ satisfies the constraint C_i. As T'' makes up at least a $.9 - 100(\epsilon + \beta)$ fraction of all the constraints, the claim follows. □

7.5 Variable gadget proof

Proof of Lemma 7.2. Because H is abelian, it is isomorphic to the group $\mathbb{Z}_{p_1} \oplus \mathbb{Z}_{p_2} \oplus \cdots \oplus \mathbb{Z}_{p_t}$, for some integers $p_1, \ldots, p_t$. Thus, we will without loss of generality assume that it equals this group. The group H can be visualized as being a p_1-by-p_2-by-…-by-p_t grid. For example, the case when $H = \mathbb{Z}_3 \oplus \mathbb{Z}_5$ is depicted in Figure 1.

The important subsets of H for our gadget construction will be those in which a specified coordinate is allowed to be free, and the rest of the coordinates are fixed. We will call these subsets *rows*. For example, a 1-row of H consists of the elements

$$(0, a_2, \ldots, a_t), (1, a_2, \ldots, a_t), \ldots, (p_t - 1, a_2, \ldots, a_t),$$

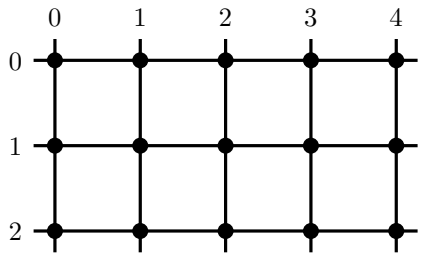

Figure 1: A grid representation of the group $\mathbb{Z}_3 \oplus \mathbb{Z}_5$. Each intersection point corresponds to a group element.

where a_i is an element of $\mathbb{Z}_{p_i}$ for each i. For a coordinate $i \in [t]$, define an i-row analogously. In Figure 1, the rows of the grid correspond to 1-rows, and the columns correspond to 2-rows.

We now construct the gadget in two steps. Each step will involve adding a set of vertices and edges which is local to each i-row.

The row gadget: Consider the group $H = \mathbb{Z}_3 \oplus \mathbb{Z}_3$. In this case, the function $f : H \to H$ in which $f(a, b) = (b, a)$ is an automorphism on H. However, it is not of the form required by the soundness condition; one way to show this is that $f(1, 0) - f(0, 0) = (0, 1)$, whereas the soundness requires that it equal $(1, 0)$. The row gadget will disallow such automorphisms by forcing any automorphism to maintain the order of coordinates.

We define the gadget G_{row} as follows: the graph initially consists of a vertex "b" for each group element $b \in H$. These are the variable vertices, the set of which we call V_x. (In the rest of the paper, we are naming these vertices "$x \mapsto b$". However, we are opting for the simpler option for this gadget construction.) Note that $V_x = H$, and so i-rows of H are i-rows of V_x, and vice versa. Now, if $t = 1$, then we stop here. Otherwise, for each coordinate i, let R be an i-row in V_x. Create a star graph with i degree-one outer nodes and one degree-i inner node, and attach this inner node to every vertex in R. Doing this for each i-row completes the construction of G_{row}. Set A_{row} to be the set of auxiliary vertices in G_{row}.

First, we show that the row gadget satisfies the necessary completeness condition:

Proposition 7.14. *For each $a \in H$, G_{row} has an automorphism f_a such that $f_a(b) = a + b$ for each $b \in H$.*

Proof. For $j \in [t]$, define $f_j(b) = b + e_j$, where $e_j = (0, \ldots, 0, 1, 0, \ldots, 0)$, with the 1 in the jth coordinate. If R is an i-row in V_x, then $f(R)$ is also an i-row, for each i. As a result, if we define f so that it maps the star graph connected to R to the star graph connected to $f_j(R)$, then f_j is an automorphism on G_{row}. This follows because each i-row is connected to a star graph with exactly i outer nodes.

For more general $a \in H$, f_a can be constructed by composing together many different f_j's. As automorphisms are closed under composition, this gives us an automorphism f_a for each $a \in H$. □

Unfortunately, G_{row} is not yet a variable gadget. However, it does satisfy the following key property:

Proposition 7.15. *Let $f : H \cup A_{row} \to H \cup A_{row}$ be an automorphism on G_{row}. Then if R is an i-row of V_x, $f(R)$ is also an i-row of V_x.*

Proof. If $t = 1$ then this is trivial, as the only vertices in the row graph are those in V_x, and V_x is the only row. Now assume that $t > 1$. Let $u \in A_{row}$ be the center of a star graph attached to the i-row R. By construction, u is connected to exactly i degree-one nodes, so $f(u)$ must have i degree-one neighbors as well. The only degree-one vertices in G_{row} are the outer nodes of the star graphs, so $f(u)$ must be the center of a star graph attached to an i-row. Because f is an automorphism, $f(R)$ must also be adjacent to $f(u)$. The only way this can happen is if $f(R)$ is the i-row adjacent to $f(u)$. Thus, $f(R)$ is an i-row if R is, and this happens for each row of V_x. □

Each star graph connected to an i-row adds $i+1 \leq 2t$ auxiliary vertices and $i+p_i \leq t+p_i$ edges. The number of i-rows is $|H|/p_i$, so this gives $2t|H|$ auxiliary vertices and $t|H| + |H| \leq 2t|H|$ edges for each coordinate i. In total, the G_{row} gadget contains $2t^2|H|$ auxiliary vertices and $2t^2|H|$ edges.

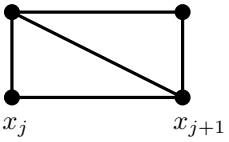

(a) The basic unit of the cycle gadget. The vertices labeled x_j and x_{j+1} are variable vertices, whereas the two unlabeled vertices are auxiliary vertices.

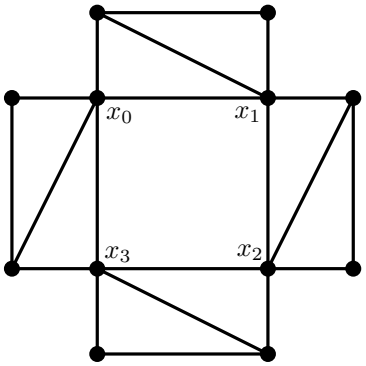

(b) The cycle gadget applied to $\mathbb{Z}_4$.

Figure 2: The cycle gadget construction. Each i-row has such a gadget placed on top of it.

The cycle gadget: Consider again the group $H = \mathbb{Z}_3 \oplus \mathbb{Z}_3$. In this case, the bijection $f : H \to H$ in which $f(a, b) = (a^3 + a, b)$ is *not* an automorphism on H. However, it can be extended to an automorphism on G_{row}. This is because f respects the order of the coordinates, and this is essentially all that the row gadget checks. The cycle gadget will enforce that each individual coordinate behaves appropriately.

Given the gadget G_{row}, we extend it to a variable gadget G as follows: pick a coordinate $i \in [t]$, and assume without loss of generality that $i = 1$. Let R be a 1-row in V_x. If $p_1 = 2$, then add a single edge between the two vertices in R. Otherwise, $p_1 > 2$, and R can be written as

$$x_0 = (0, a_2, \ldots, a_t), \quad x_1 = (1, a_2, \ldots, a_t), \quad \ldots, \quad x_{p_1-1} = (p_1 - 1, a_2, \ldots, a_t).$$

In this case, we imagine $x_0, x_1, \ldots, x_{p_1-1}$ as being arranged clockwise around a circle, and between the elements x_j and $x_j + 1$ (where addition is modulo p_1), we add the vertices and edges shown in Figure 2a. For example, Figure 2b shows the result when $p_1 = 4$. Doing this for each i-row and adding the result on top of G_{row} completes the construction of G. Set A_{cycle} to be the set of additional auxiliary vertices created in this step, which is disjoint from A_{row}.

We first show that G satisfies the completeness condition:

Proposition 7.16. *For each $a \in H$, G has an automorphism f_a such that $f_a(b) = a + b$ for each $b \in H$.*

Proof. Define f_j on the variable vertices and the star graphs as in Proposition 7.14. We need to define f_j on the cycle gadgets as well, and the same reasoning as before shows that because f_j maps i-rows to i-rows (and it maintains the relative order of the vertices in each rows), we can set f_j to locally map the cycle gadget from row R to row $f_j(R)$ so that f_j is an automorphism on G. And as before, it is easy to extend this to hold for all f_a. □

In addition, it satisfies the soundness condition:

Proposition 7.17. *Let f be an automorphism on G. Then there exists an element $a \in H$ such that $f(b) = a + b$ for each $b \in H$.*

Proof. Let f be an automorphism on G, and let R be an i-row of V_x. Note that none of the vertices in A_{cycle} have degree one. Therefore, the only degree-one vertices in G are in the star graphs from the row gadget. This is enough to reuse the argument in Proposition 7.15 to show that $f(V_x) = V_x$, $f(A_{row}) = A_{row}$, and $f(A_{cycle}) = A_{cycle}$. In addition, we may apply Proposition 7.15 to R, which shows that $f(R)$ is also an i-row of V_x.

Write the elements of R, as in the above paragraph, as $x_0, \ldots, x_{p_i-1}$, where x_j has a j in the ith coordinate. We have the relationship $x_{j+1} = x_j + e_i$. Let u be the vertex adjacent to both x_j and x_{j+1} in Figure 2a, and let v be the vertex adjacent to both u and x_{j+1}. The only row that x_j and x_{j+1} are both in is R. This means that u is the *only* vertex in A_{cycle} adjacent to both x_j and x_{j+1}, and likewise v is the only vertex in A_{cycle} adjacent to both u and x_{j+1}. This means that if f is an automorphism, $f(u)$ must be the only vertex in A_{cycle} adjacent to both

$f(x_j)$ and $f(x_{j+1})$, and $f(v)$ must be the only vertex in A_{cycle} adjacent to both $f(u)$ and $f(x_{j+1})$. The only way this is possible is if $f(x_{j+1}) = f(x_j) + e_i$. Rewriting x_{j+1} as $x_j + e_i$, we conclude that $f(x_j + e_i) = f(x_j) + e_i$.

The above argument only used that x_j and x_{j+1} differed by one in the ith coordinate. As a result, we know that $f(b + e_i) = f(b) + e_i$ for any coordinate i. More generally, we know that $f(a + b) = f(a) + b$, which can be shown by applying the previous fact coordinate-by-coordinate. Set $a := f(0)$. Then $f(b) = f(0) + b = a + b$. This completes the soundness case. □

Each time the cyclic gadget is applied to an i-row, it adds $2p_i$ auxiliary vertices and $5p_i$ edges to the graph. As the number of i-rows is $|H|/p_i$, this adds $2|H|$ vertices and $5|H|$ edges per coordinate i, for a total of $2t|H|$ vertices and $5t|H|$ edges. Combining this with the row gadget, the vertex gadget creates at most $4t^2|H|$ vertices and $7t^2|H|$ edges. Using the fact that $p_1 \cdots p_t = |H|$, we see that $t \leq \log_2 |H|$, giving the bound in the lemma statement.

As for the size of the maximum clique, we consider three cases. If $|H| = 2$, then the variable gadget is a 2-clique. If $|H| = 3$, then the variable gadget is as in Figure 2b, except with a triangle in the center rather than a square. It can be checked that the largest clique in this case is of size three. Finally, if $|H| > 3$, any clique containing an auxiliary variable from a cycle gadget must be of size at most three, as Figure 2a verifies. Next, any group of variable vertices can appear in a clique together only if they appear in the same i-row together. Furthermore, the largest clique in an i-row is of size two unless $p_i = 3$, meaning the three variable vertices in that i-row form a triangle. In this case (supposing $t > 1$), they will also all be connected to the center of a star graph by the row gadget. This is the only way to form a clique of size four, and no larger clique is possible. Thus, the maximum clique size is at most $\min\{4, |H|\}$. □

7.6 Label-extended graph proof

Proof of Lemma 7.3. We handle the completeness and soundness cases separately:

Completeness. Given α, define f as follows: first, $f(x_i \mapsto b) := x_i \mapsto (b - \alpha(x_i))$ for each $i \in [k]$ and $b \in H$. Next, for each assignment β which satisfies C, set $f(\beta) := \beta - \alpha$, using coordinate-wise subtraction. For this definition to be legitimate, we must ensure that $\beta - \alpha$ is a satisfying assignment of ψ. Because C is an $\mathsf{Additive{-}CSP}(\psi)$ constraint, it can be written as

$$\psi(x_1 + a_1, \ldots, x_k + a_k) = 1.$$

Because α and β both satisfy C, they can be written as $\alpha' - (a_1, \ldots, a_k)$ and $\beta' - (a_1, \ldots, a_k)$, where α' and β' are satisfying assignments of φ. Then by the subgroup property of φ, $\beta - \alpha = \beta' - \alpha' \in \varphi$. Therefore, f is a legitimate bijection between the vertices of G and $\underline{G}$.

Given an assignment β, there is an edge in G between β and $x_i \mapsto \beta(i)$, for each $i \in [k]$. Then $f(\beta) = \beta - \alpha$, and $f(x_i \mapsto \beta(i)) = x_i \mapsto (\beta(i) - \alpha(i))$, both of which are connected by an edge in $\underline{G}$. Thus, f is a homomorphism.

Soundness. Let β be the assignment in V_C which f maps to the all-0's assignment. Then for each $i \in [k]$, $x_i \mapsto \beta(i)$ is connected to β in G, so $f(x_i \mapsto \beta(i)) = x_i \mapsto (\beta(i) - \alpha(i))$ must be connected to 0 in $\underline{G}$. This means that $\beta(i) - \alpha(i) = 0$ for all $i \in [k]$. In particular, $\alpha = \beta$, and because β is a vertex in V_C, it (and by extension α) is a satisfying assignment of C. □

8 Conclusions

We have shown SOS/Lasserre gaps for GISO and hardness of ROBUSTGISO based on various average case hardness assumptions. The gaps we obtained are tiny. Although we did not attempt to optimize over the parameters in order to maximize the gaps, we believe the current approach cannot make the soundness as small as an arbitrary constant. Therefore one intriguing problem at this point is to find a way amplifying the gaps. A possible approach is to prove a "parallel repetition theorem for GISO", i.e. find a graph product operator $\square$ such that $\mathsf{GI}(G^{\square\ell}, H^{\square\ell}) \leq \mathsf{GI}(G, H)^{\omega_\ell(1)}$, where $\omega_\ell(1)$ means a function which goes to $+\infty$ as $\ell \to +\infty$. A natural

candidate for $\square$ is the tensor product $\otimes$. It is known that ([McK71]), for (exact) GRAPHISOMORPHISM, if both G and H are connected and non-bipartite graphs, $G \otimes G$ is not isomorphic to $H \otimes H$ when G is not isomorphic to H. But how the tensor product works with approximate GRAPHISOMORPHISM remains to be explored.

Another future direction is to prove robust asymmetry property for random d-regular graphs. Since in this model we do not have the problem of many isolated vertices, even for constant d, we might be able to prove a $(\beta, \Omega(\beta))$-asymmetry property for beta being as small as $\Theta\left(\frac{1}{n}\right)$, where in contrast, there is an $\exp(-O(c))$ lower bound for β in the $\mathcal{G}_{n,cn}$ model.

References

[ABS10] Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for Unique Games and related problems. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 563–572, 2010. 1

[AFK02] Sanjeev Arora, Alan Frieze, and Haim Kaplan. A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. *Mathematical Programming, Series A*, 92(1):1–36, 2002. 4

[AK05] V. Arvind and Johannes Köbler. Isomorphism testing: Perspective and open problems. *Bulletin of the European Association for Theoretical Computer Science*, 86:66–84, 2005. 1

[AKKV12] Vikraman Arvind, Johannes Köbler, Sebastian Kuhnert, and Yadu Vasudev. Approximate graph isomorphism. In *Proceedings of the 37th Annual Innovations in Theoretical Computer Science conference*, pages 100–111, 2012. 4

[Ale03] Michael Alekhnovich. More on average case vs. approximation complexity. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 298–307, 2003. 4

[AM13] Albert Atserias and Elitza Maneva. Sherali–Adams relaxations and indistinguishability in counting logics. *SIAM Journal on Computing*, 42(1):112–137, 2013. 2

[AR05] Dorit Aharonov and Oded Regev. Lattice problems in $\mathrm{NP} \cap \mathrm{coNP}$. *Journal of the ACM*, 52(5):749–765, 2005. 1

[Bab85] László Babai. Trading group theory for randomness. In *Proceedings of the 17th Annual ACM Symposium on Theory of Computing*, pages 421–429, 1985. 1

[Bab95] László Babai. Automorphism groups, isomorphism, reconstruction. In *Handbook of combinatorics, Vol. 1, 2*, pages 1447–1540. Elsevier, Amsterdam, 1995. 1

[BBH+12] Boaz Barak, Fernando Brandão, Aram Harrow, Jonathan Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 307–326, 2012. 2, 6

[BES80] László Babai, Paul Erdős, and Stanley Selkow. Random graph isomorphism. *SIAM Journal on Computing*, 9(3):628–635, 1980. 2

[BHZ87] Ravi Boppana, Johan Håstad, and Stathis Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25(2):127–132, 1987. 1

[BK79] László Babai and Ludik Kučera. Canonical labelling of graphs in linear average time. In *Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science*, pages 39–46, 1979. 2

[BK12] Libor Barto and Marcin Kozik. Robust satisfiability of constraint satisfaction problems. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 931–940, 2012. 3

[BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM*, 50(4):506–519, 2003. 4

[BL83] László Babai and Eugene Luks. Canonical labeling of graphs. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing*, pages 171–183, 1983. 1, 2

[Bol82] Béla Bollobás. The asymptotic number of unlabelled regular graphs. *The Journal of the London Mathematical Society, Second Series*, 26(2):201–206, 1982. 4

[CFI92] Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992. 0, 1, 2, 3, 7

[Cha13] Siu On Chan. Approximation resistance from pairwise independent subgroups. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, pages 447–456. ACM, 2013. 4, 18

[Cod11] Paolo Codenotti. *Testing isomorphism of combinatorial and algebraic structures*. PhD thesis, University of Chicago, 2011. 1

[CSS14] Paolo Codenotti, Grant Schoenebeck, and Aaron Snook. Graph isomorphism and the Lasserre hierarchy. Technical report, arXiv:1401.0758, 2014. 3

[DK12] Victor Dalmau and Andrei Krokhin. Robust satisfiability for CSPs: Algorithmic and hardness results. Unpublished, 2012. 3

[ER63] Paul Erdős and Alfréd Rényi. Asymmetric graphs. *Acta Mathematica Hungarica*, 14(3):295–315, 1963. 3, 4, 5

[Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 543–543, 2002. 0, 1, 3

[GG00] Oded Goldreich and Shafi Goldwasser. On the limits of nonapproximability of lattice problems. *Journal of Computer and System Sciences*, 60(3):540–563, 2000. 1

[GO12] Martin Grohe and Martin Otto. Pebble games and linear equations. Technical Report 1204.1990, arXiv, 2012. 2

[Gri01] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1-2):613–622, 2001. 3

[Gro12] Martin Grohe. Fixed-point definability and polynomial time on graphs with excluded minors. *Journal of the ACM*, 59(5):Article 27, 2012. 2

[GZ11] Venkatesan Guruswami and Yuan Zhou. Tight bounds on the approximability of almost-satisfiable Horn SAT and Exact Hitting Set. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1574–1589, 2011. 3

[Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001. 4

[Imm81] Neil Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22(3):384–406, 1981. 2

[JP83] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, 11(1):286–295, 1983. 13

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 767–775, 2002. 1

[KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for Max-Cut and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007. 4

[Köb06] Johannes Köbler. On graph isomorphism for restricted graph classes. In Arnold Breckmann, Ulrich Berger, Benedikt Löwe, and John Tucker, editors, *Logical Approaches to Computational Barriers*, number 3988 in LNCS, pages 241–256. Springer-Verlag, 2006. 1

[KOT+12] Gábor Kun, Ryan O'Donnell, Suguru Tamaki, Yuichi Yoshida, and Yuan Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *Proceedings of the 3nd Annual Innovations in Theoretical Computer Science conference*, pages 484–495, 2012. 3

[KOTZ13] Manuel Kauers, Ryan O'Donnell, Li-Yang Tan, and Yuan Zhou. Hypercontractive inequalities via SOS, with an application to Vertex-Cover. Manuscript, 2013. 2

[KST93] Johannes Köbler, Uwe Schöning, and Jacobo Torán. *The graph isomorphism problem: its structural complexity*. Progress in Theoretical Computer Science. Birkhäuser Boston Inc., Boston, MA, 1993. 1

[KSV02] Jeong Han Kim, Benny Sudakov, and Van Vu. On the asymmetry of random regular graphs and random graphs. *Random Structures & Algorithms*, 21(3-4):216–224, 2002. 4

[Las00] Jean Lasserre. Optimisation globale et théorie des moments. *Comptes Rendus de l'Académie des Sciences*, 331(11):929–934, 2000. 2

[McK71] Ralph McKenzie. Cardinal multiplication of structures with a reflexive relation. *Fundamenta Mathematicae*, 70(1):59–101, 1971. 26

[McK81] Brendan McKay. Practical graph isomorphism. *Congressus Numerantium*, 30:45–87, 1981. 1

[MOO10] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *Annals of Mathematics*, 171(1), 2010. 4

[MW84] Brendan McKay and Nicholas C. Wormald. Automorphisms of random graphs with specified vertices. *Combinatorica*, 4(4):325–338, 1984. 4

[OZ13] Ryan O'Donnell and Yuan Zhou. Approximability and proof complexity. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1537–1556, 2013. 2

[Par00] Pablo Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, 2000. 2

[RC77] Ronald Read and Derek Corneil. The graph isomorphism disease. *Journal of Graph Theory*, 1(4):339–363, 1977. 1

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 84–93, 2005. 4

[RSU94] Motakuri Ramana, Edward Scheinerman, and Daniel Ullman. Fractional isomorphism of graphs. *Discrete Mathematics*, 132(1-3):247–265, 1994. 2

[SA90] Hanif Sherali and Warren Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990. 2

[Sch88] Uwe Schöning. Graph isomorphism is in the low hierarchy. *Journal of Computer and System Sciences*, 37(3):312–323, 1988. 1

[Sch08] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 593–602, 2008. 3, 7

[Spi96] Daniel Spielman. Faster isomorphism testing of strongly regular graphs. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, pages 576–584, 1996. 2

[Tor04] Jacobo Torán. On the hardness of graph isomorphism. *SIAM Journal on Computing*, 33(5):1093–1108, 2004. 1, 2

[TSSW00] Luca Trevisan, Gregory Sorkin, Madhu Sudan, and David Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000. 4

[Tul09] Madhur Tulsiani. CSP gaps and reductions in the Lasserre hierarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, pages 303–312, 2009. 3, 18

[Wei76] Boris Weisfeiler, editor. *On construction and identification of graphs*. Lecture Notes in Mathematics, Vol. 558. Springer–Verlag, Berlin, 1976. 1, 2

[WYZV13] Yu Wu, Yuichi Yoshida, Yuan Zhou, and Aravindan Vijayaraghavan. Graph isomorphism: approximate and robust. Manuscript, 2013. 3

[Zwi98] Uri Zwick. Finding almost-satisfying assignments. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 551–560, 1998. 3